NA-ASC-114R-07-Vol. 2, Rev. 0.5
UCRL-TR-234798

FY08-09 Implementation Plan Volume 2, Rev. 0.5

September 20, 2007

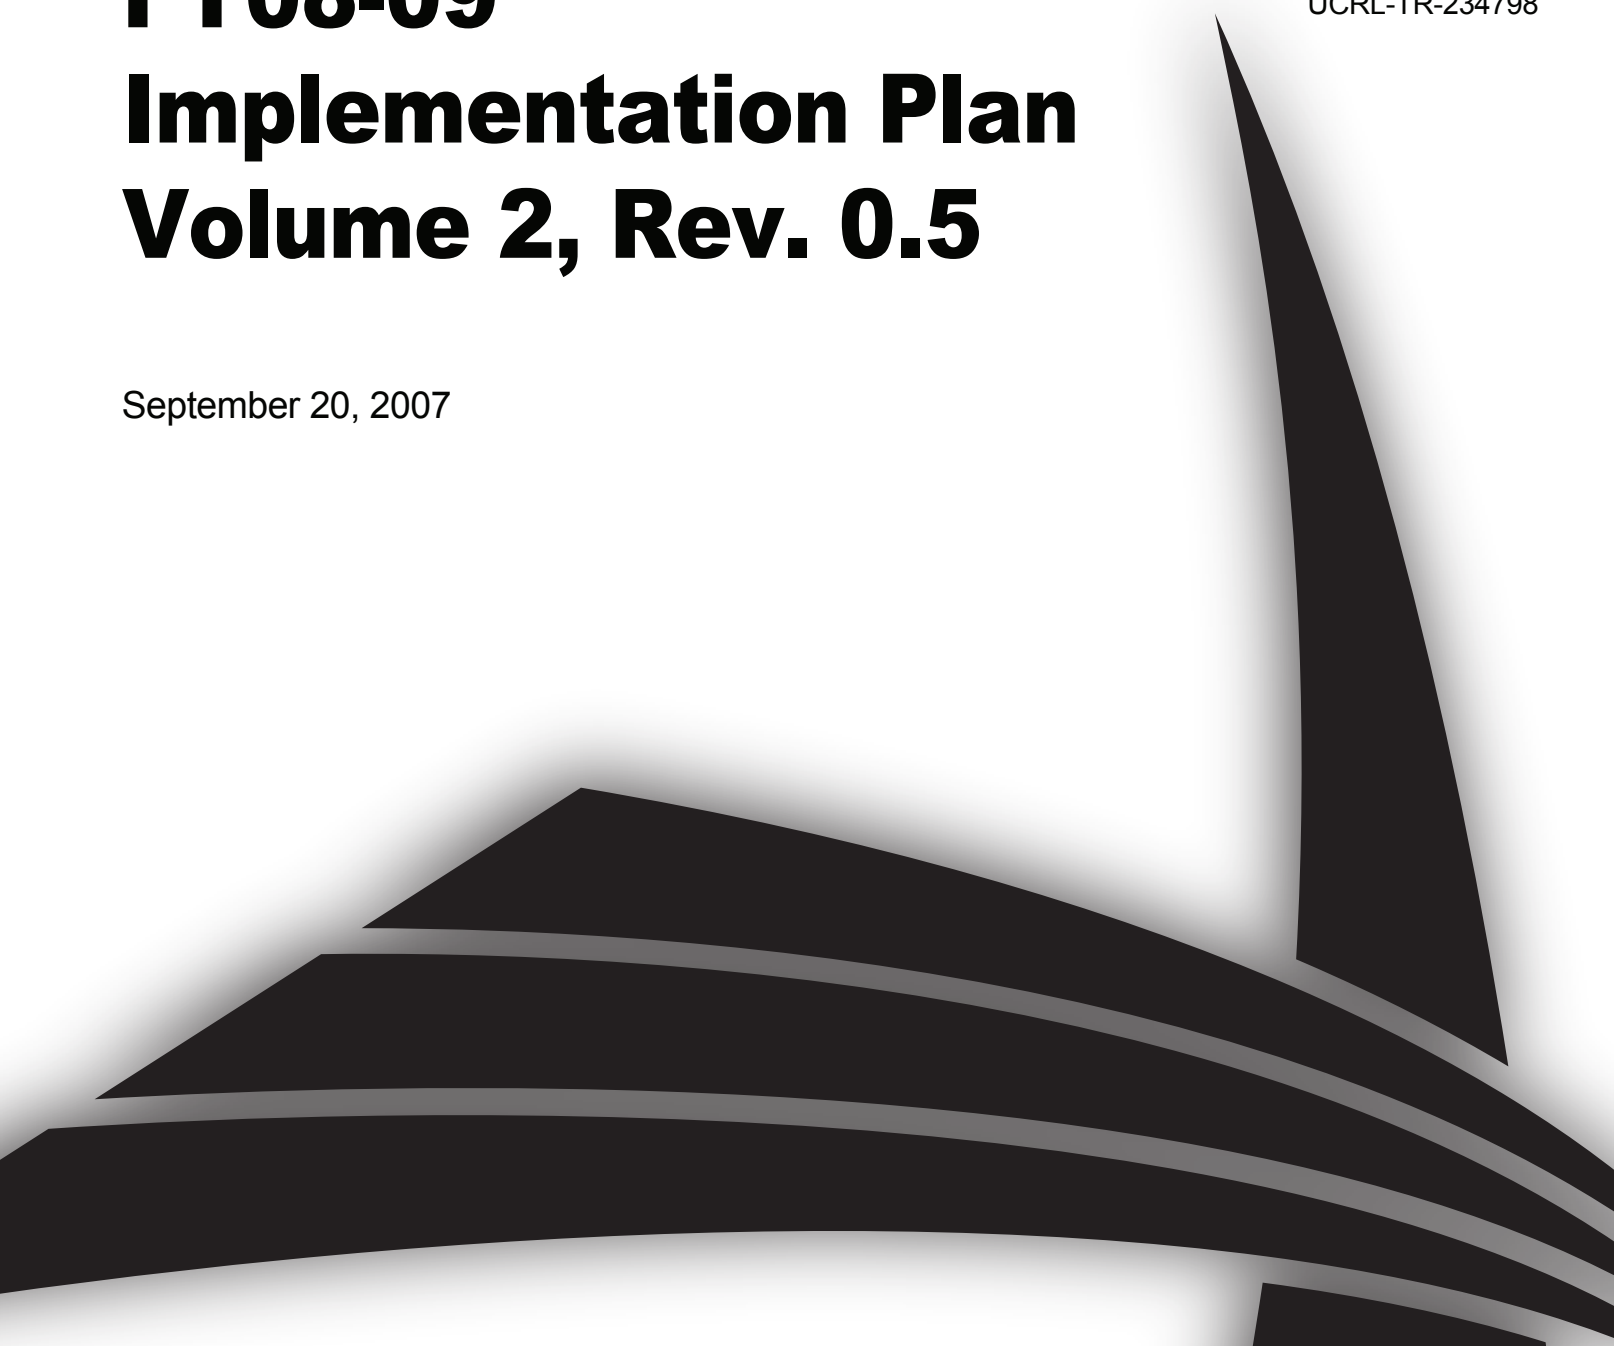

This work was performed under the auspices of the U.S. Department of Energy by University of California, Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

Advanced Simulation and Computing

FY08–09 IMPLEMENTATION PLAN
Volume 2, Rev. 0.5

September 20, 2007

Approved by:

Dimitri Kusnezov, NNSA ASC Program Director	Signature	*9/20/07* Date
Thomas Bickel, SNL ASC Executive	Signature	*9/20/07* Date
Michel McCoy, LLNL ASC Executive	Signature	*9/20/07* Date
John Hopson, LANL ASC Executive	Signature	*9/20/07* Date

ASC Focal Point
Dimitri Kusnezov
NA 114
Tele.: 202-586-1800
FAX: 202-586-0405
dimitri.kusnezov@nnsa.doe.gov

IP Focal Point
Njema Frazier
NA 114
Tele.: 202-586-5789
FAX: 202-586-7754
njema.frazier@nnsa.doe.gov

Implementation Plan Contents at a Glance

Section No./Title		Vol. 1	Vol. 2	Vol. 3
I.	Executive Summary			
II.	Introduction			
III.	Accomplishments			
IV.	Product Descriptions			
V.	ASC Level 1 and 2 Milestones			
VI.	ASC Roadmap Drivers for FY08–FY09			
VII.	ASC Risk Management			
VIII.	Performance Measures			
IX.	Budget			
Appendix A.	Glossary			
Appendix B.	Points of Contact			
Appendix C.	1.5.1.4.12 Alliance Centers			
Appendix D.	ASC Obligation/Cost Plan			
Appendix E.	ASC Performance Measurement Data for FY08			

Contents

List of Tables

List of Figures

I. Executive Summary

The Stockpile Stewardship Program (SSP) is a single, highly integrated technical program for maintaining the surety and reliability of the U.S. nuclear stockpile. The SSP uses past nuclear test data along with current and future non-nuclear test data, computational modeling and simulation, and experimental facilities to advance understanding of nuclear weapons. It includes stockpile surveillance, experimental research, development and engineering programs, and an appropriately scaled production capability to support stockpile requirements. This integrated national program requires the continued use of current facilities and programs along with new experimental facilities and computational enhancements to support these programs.

The Advanced Simulation and Computing Program (ASC)[1] is a cornerstone of the SSP, providing simulation capabilities and computational resources to support the annual stockpile assessment and certification, to study advanced nuclear-weapons design and manufacturing processes, to analyze accident scenarios and weapons aging, and to provide the tools to enable Stockpile Life Extension Programs (SLEPs) and the resolution of Significant Finding Investigations (SFIs). This requires a balanced resource, including technical staff, hardware, simulation software, and computer science solutions.

In its first decade, the ASC strategy focused on demonstrating simulation capabilities of unprecedented scale in three spatial dimensions. In its second decade, ASC is focused on increasing its predictive capabilities in a three-dimensional simulation environment while maintaining the support to the SSP. The program continues to improve its unique tools for solving progressively more difficult stockpile problems (focused on sufficient resolution, dimensionality and scientific details); to quantify critical margins and uncertainties (QMU); and to resolve increasingly difficult analyses needed for the SSP. Moreover, ASC has restructured its business model from one that was very successful in delivering an initial capability to one that is integrated and focused on requirements-driven products that address long-standing technical questions related to enhanced predictive capability in the simulation tools.

ASC must continue to meet three objectives:

- **Objective 1. Robust Tools.** Develop robust models, codes, and computational techniques to support stockpile needs such as refurbishments, SFIs, LEPs, annual assessments, and evolving future requirements.
- **Objective 2. Prediction through Simulation.** Deliver validated physics and engineering tools to enable simulations of nuclear-weapons performances in a variety of operational environments and physical regimes and to enable risk-informed decisions about the performance, safety, and reliability of the stockpile.
- **Objective 3. Balanced Operational Infrastructure.** Implement a balanced computing platform acquisition strategy and operational infrastructure to meet Directed Stockpile Work (DSW) and SSP needs for capacity and high-end simulation capabilities.

[1] In FY02 the Advanced Simulation and Computing (ASC) Program evolved from the Accelerated Strategic Computing Initiative (ASCI).

II. Introduction

The ASC Program supports the National Nuclear Security Administration's (NNSA's) long-term strategic goal of Nuclear Weapons Stewardship: *"ensure that our nuclear weapons continue to serve their essential deterrence role by maintaining and enhancing the safety, security, and reliability of the U.S. nuclear weapons stockpile."*[2]

In 1996, ASCI—the Accelerated Strategic Computing Initiative—was established as an essential element of the SSP to provide nuclear weapons simulation and modeling capabilities.

In 2000, the NNSA was established to carry out the national security responsibilities of the Department of Energy, including maintenance of a safe, secure, and reliable stockpile of nuclear weapons and associated materials capabilities and technologies.

Shortly thereafter, in 2002, ASCI matured from an initiative to a recognized program and was renamed the Advanced Simulation and Computing (ASC) Program.

Prior to the start of the nuclear testing moratorium in October 1992, the nuclear weapons stockpile was maintained through (1) underground nuclear testing and surveillance activities and (2) "modernization" (i.e., development of new weapons systems). A consequence of the nuclear test ban is that the safety, performance, and reliability of U.S. nuclear weapons must be ensured by other means for systems far beyond the lifetimes originally envisioned when the weapons were designed.

NNSA will carry out its responsibilities through the twenty-first century in accordance with the current Administration's vision and the Nuclear Posture Review (NPR) guidance. NNSA Administrator Ambassador Brooks summarized[3] the NNSA objectives for SSP as follows:

"We will continue to lead the way to a safer world through the deep reductions in nuclear forces codified by the Moscow Treaty, through Nunn-Lugar and other cooperative threat reduction efforts, and through other actions. At the same time, although conventional forces will assume a larger share of the deterrent role, we will maintain an effective, reliable, and capable—though smaller—nuclear force as a hedge against a future that is uncertain and in a world in which substantial nuclear arsenals remain. Our ongoing efforts to reduce the current stockpile to the minimum consistent with national security requirements, to address options for transformation of this smaller stockpile, and to create a responsive nuclear weapons infrastructure are key elements of the Administration's national security strategy…"

A truly responsive infrastructure will allow us to address and resolve any stockpile problems uncovered in our surveillance program; to adapt weapons (achieve a capability to modify or repackage existing warheads within 18 months of a decision to enter engineering development); to be able to design, develop, and initially produce a new warhead within three to four years of a decision to do so;[4] to restore production

[2] NNSA Strategic Plan, page 8.

[3] Speech presented to the Heritage Foundation Conference: *U.S. Strategic Command: Beyond the War on Terrorism*, May 12, 2004.

[4] While there are no plans to develop new weapons, gaining the capability is an important prerequisite to deep reductions in the nuclear stockpile.

capacity to produce new warheads in sufficient quantities to meet any defense needs that arise without disrupting ongoing refurbishments; to ensure that services such as warhead transportation, tritium support, and other ongoing support efforts are capable of being carried out on a time scale consistent with the Department of Defense's ability to deploy weapons; and to improve test readiness (an 18-month test readiness posture) in order to be able to diagnose a problem and design a test that could confirm the problem or certify the solution (without assuming any resumption of nuclear testing).

Additionally, the NPR guidance has directed that NNSA maintain a research and development and manufacturing base that ensures the long-term effectiveness of the nation's stockpile and begin a modest effort to examine concepts (for example, Advanced Concepts Initiatives) that could be deployed to further enhance the deterrent capabilities of the stockpile in response to the national security challenges of the twenty-first century.

The ASC Program plays a vital role in the NNSA infrastructure and its ability to respond to the NPR guidance. The program focuses on development of modern simulation tools that can provide insights into stockpile problems, provide tools with which designers and analysts can certify nuclear weapons, and guide any necessary modifications in nuclear warheads and the underpinning manufacturing processes. Additionally, ASC is enhancing the predictive capability necessary to evaluate weapons effects, design experiments, and ensure test readiness.

ASC continues to improve its unique tools to solve progressively more difficult stockpile problems, with a focus on sufficient resolution, dimensionality, and scientific details, to enable QMU and to resolve the increasingly difficult analyses needed for stockpile stewardship. The DSW provides requirements for simulation, including planned SLEPs, stockpile support activities that may be ongoing or require short-term urgent response, and requirements for future capabilities to meet longer-term stockpile needs, such as the Reliable Replacement Warhead (RRW). Thus, ASC's advancing leading-edge technology in high-performance computing and predictive simulation meets these short- and long-term needs, including the annual assessments and certifications and SFIs. The following section lists past, present, and planned ASC contributions to meet these needs.

ASC Contributions to the Stockpile Stewardship Program

In FY96, ASCI Red was delivered. Red, the world's first teraFLOPS supercomputer, was upgraded to more than 3 teraFLOPS in FY99 and was retired from service in September 2005.

In FY98, ASCI Blue Pacific and ASCI Blue Mountain were delivered. These platforms were the first 3-teraops systems in the world and have both since been decommissioned.

In FY00, ASCI successfully demonstrated the first-ever three dimensional (3D) simulation of a nuclear weapon primary explosion and the visualization capability to analyze the results; ASCI successfully demonstrated the first-ever 3D hostile-environment simulation; and ASCI accepted delivery of ASCI White, a 12.3-teraops supercomputer, which has since been retired from service.

In FY01, ASCI successfully demonstrated simulation of a 3D nuclear weapon secondary explosion; ASCI delivered a fully functional Problem Solving Environment for ASCI White; ASCI demonstrated high-bandwidth distance computing between the three national laboratories; and ASCI demonstrated the initial validation methodology for early primary behavior. Lastly, ASCI completed the 3D analysis for a stockpile-to-target sequence (STS) for normal environments.

In FY02, ASCI demonstrated 3D system simulation of a full-system (primary and secondary) thermonuclear weapon explosion, and ASCI completed the 3D analysis for an STS abnormal-environment crash-and-burn accident involving a nuclear weapon.

In FY03, ASCI delivered a nuclear safety simulation of a complex, abnormal, explosive initiation scenario; ASCI demonstrated the capability of computing electrical responses of a weapons system in a hostile (nuclear) environment; and ASCI delivered an operational 20-teraops platform on the ASCI Q machine, which currently is being retired from service.

In FY04, ASC provided simulation codes with focused model validation to support the annual certification of the stockpile and to assess manufacturing options. ASC supported the life-extension refurbishments of the W76 and W80, in addition to the W88 pit certification. In addition, ASC provided the simulation capabilities to design various non-nuclear experiments and diagnostics.

In FY05, ASC identified and documented SSP requirements to move beyond a 100-teraops computing platform to a petaFLOPS-class system; ASC delivered a metallurgical structural model for aging to support pit-lifetime estimations, including spiked-plutonium alloy. In addition, ASC provided the necessary simulation codes to support test readiness as part of NNSA's national priorities.

In FY06, ASC delivered the capability to perform nuclear performance simulations and engineering simulations related to the W76/W80 LEPs to assess performance over relevant operational ranges, with assessments of uncertainty levels for selected sets of simulations. The deliverables of this milestone were demonstrated through 2D and 3D physics and engineering simulations. The engineering simulations analyzed system behavior in abnormal thermal environments and mechanical response of systems to hostile blasts. Additionally, confidence measures and methods for uncertainty quantification (UQ) were developed to support weapons certification and QMU Level 1 milestones.

In FY07, ASC supported the completion of the W76-1 and W88 warhead certification, using quantified design margins and uncertainties; ASC also provided two robust 100-teraFLOPS-platform production environments by IBM and CRAY, supporting DSW and Campaign simulation requirements, respectively. One of the original ASCI program Level 1 milestones was completed when the ASC Purple system was formally declared "general available." This was augmented by the 360-teraFLOPS ASC BlueGene/L system, which provided additional capability for science campaigns. The ASC-funded partnerships between SNL/Cray and LLNL/IBM have transformed the supercomputer industry. There are 34 "Blue Gene Solution" systems on the Top 500 list and 35 Cray sales based on the Red Storm architecture.

By FY08, ASC will deliver the codes for experiment and diagnostic design to support the CD-4 approval on the National Ignition Facility (NIF). An advanced architecture platform capable of sustaining a 1-petaFLOPS benchmark will be sited at LANL.

By FY09, a modern baseline of all enduring stockpile systems, using ASC codes, will be completed.

In FY10 and beyond, ASC will continue to deliver codes for experiment and diagnostic design to support the indirect-drive ignition experiments on the NIF and will continue to improve confidence and response time for predictive capabilities to answer questions of vital importance to the SSP.

Table II-1. Defense Program Campaigns

Campaign Number	Campaign Title
C1	Primary Assessment Technology and Test Readiness
C2	Dynamic Materials Properties
C3	Advanced Radiography
C4	Secondary Assessment Technology
C5	Enhanced Surety
C6	Weapon Systems Engineering Assessment Technology
C7	Nuclear Survivability
C8	Enhanced Surveillance
C9	Advanced Design & Production Technologies
C10	Inertial Confinement Fusion Ignition and High Yield Campaign
C11	Advanced Simulation and Computing
C12	Pit Manufacturing and Certification Campaign
C15	Non-Nuclear Readiness
C16	Materials Readiness
C18	Engineering Campaigns Construction Activities
C19	Advanced Design & Production Technologies Readiness

The National Work Breakdown Structure

ASC's program structure is based on the new national work breakdown structure (nWBS), described in the ASC Business Model (NA-ASC-104R-05-Vol.1-Rev.5).

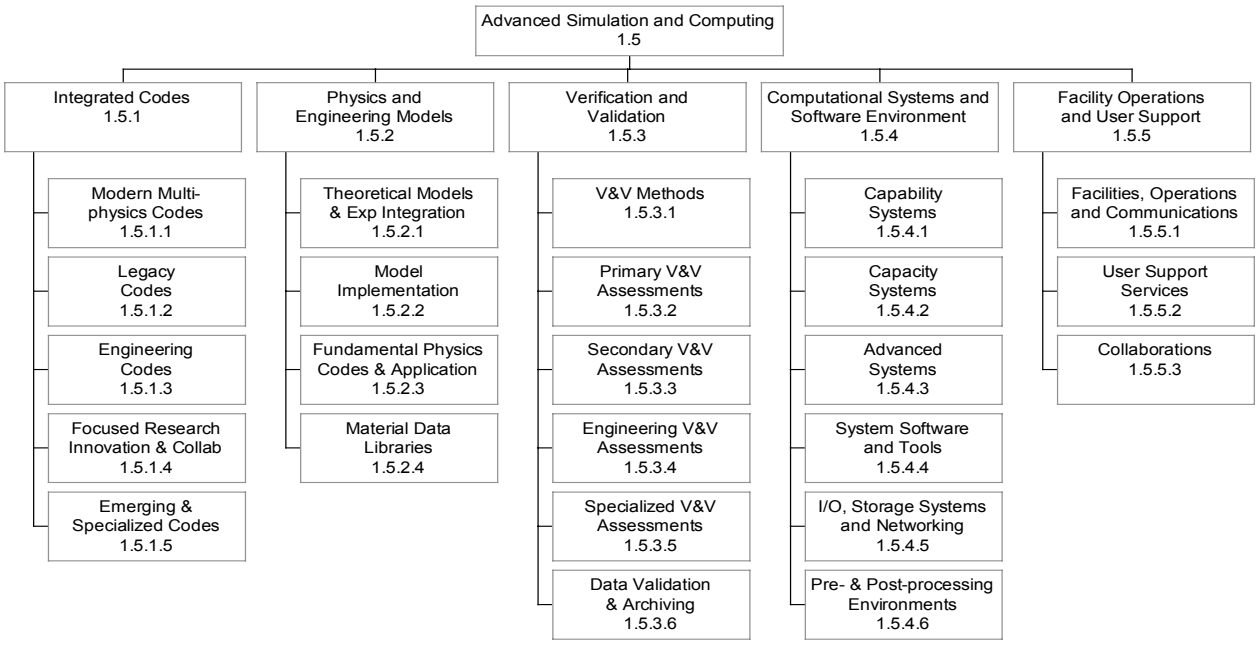

Figure II-1. The ASC program structure is based on the new national work breakdown structure.

Sub-Programs

As the chart shows, ASC is divided into five sub-programs:

- Integrated Codes
- Physics and Engineering Models
- Verification and Validation
- Computational Systems and Software Environment
- Facility Operations and User Support

The first three sub-programs focus on improved models in the modern codes, delivery of validated tools, and response to SSP issues (for example, SFIs, LEPs, annual assessments). Key drivers are to improve the confidence in prediction through simulations; to calculate, measure, and understand the uncertainty in the predictions; and to ensure rapid delivery of simulation capabilities to the SSP.

The fourth sub-program, Computational Systems and Software Environment, ensures the development and deployment of a computing environment needed for all ASC-deployed platforms: capability, capacity, and advanced systems.[5] Not only is this sub-

[5] The ASC Program is in transition for current platforms. Future platforms will follow the Capital Acquisition Management process identified in the NA-10 *Program Management Manual*.

program responsible for related research and technology development, but it is also responsible for planning, procurement, and quality control activities.

The fifth, and last, sub-program, Facility Operations and User Support, provides operational support for production computing and storage, user support services, and collaborative research opportunities with educational institutions, as well as programmatic support across the ASC program.

Product Deliverables

The Product deliverables are described at level 4 and span the full-scope of the program in the context of the nWBS. They describe what the Laboratories expect to provide to a given Product as a result of their activities.

Deliverables can, but do not necessarily, contribute to level 2 milestones chosen in a given fiscal year. Deliverables that do directly contribute to level 2 milestones for the fiscal year will be reviewed in the context of level 2 milestone reviews.

III. Accomplishments for FY06–FY07

ASC accomplishments from Quarter 4, fiscal year 2006, through quarter 3, fiscal year 2007, are reflected below for the Computational Systems and Software Environment (CSSE) and Facility Operations and User Support (FOUS) sub-programs.

Headquarters is pleased to highlight the outstanding achievements of the Defense Programs Contractors.

Computational Systems and Software Environment

LLNL Accomplishments for Computational Systems and Software Environment

- Deployed the ASC Purple 100 teraFLOPS system for classified general availability (GA) to the tri-lab SSP program and certified the Level 1 Milestone #350 as completed. Activities associated with this major effort cut across most CSSE and FOUS products and included contributions from SNL and LANL in testing for the milestone completion.
- Initiated tri-lab Linux Capacity Cluster 2007 (TLCC07) procurement with the request for proposal released in July 2007.
- Signed the multi-year research and development BlueGene/P/Q contract with IBM in partnership with the Office of Science and NNSA. Contract deliverables in FY07 included BlueGene/P systems packaging design, various system functionality designs, and initial hardware demonstrations.
- Submitted the Critical Decision Mission Need Package (CD-0) for the proposed Sequoia computer system at the Lawrence Livermore National Laboratory (LLNL) to the NNSA Federal Acquisition Executive, who approved it.
- Deployed and ran high performance storage system (HPSS) R6.2 at scale. LLNL was the first production site to meet this goal.
- Contributed to LLNL's ability to understand materials aging processes through performance analysis and optimization of several key applications, including participation on the 2006 Gordon Bell Prize winning Qbox team.
- Developed a strategy for large-scale debugging to ensure LLNL applications can run correctly on petascale platforms. Developed new methodologies for performance analysis and optimization (recognition as Software Track Best Paper Award at the 2007 International Parallel and Distributed Processing Symposium).
- Received ten scalable units (SUs) of the Minos, Rhea, and Hopi systems for capacity computing. Rhea and Hopi are now in production.
- Deployed Lustre parallel file system on InfiniBand across capacity clusters and major increases of Lustre storage space for classified network (to 2 petabytes) and unclassified network (to 1 petabyte).
- Procured and integrated the Prism visualization cluster into the Open Computing Facility (OCF).

LANL Accomplishments for Computational Systems and Software Environment

- The Roadrunner base system (Redtail) passed acceptance testing in December 2006 and integrated into LANL classified computing environment May 2006. An ASC Level 2 milestone for system deployment completed June 2007. Crestone and other stockpile weapons code applications were working on the Roadrunner Base system by July 2007.
- The Physics-Based Simulation Analysis project identified in high-resolution simulations two new phenomena whose importance in weapons physics was not previously understood. These discoveries suggest a new paradigm for predicting weapons performance.
- The Performance and Architecture Lab is now contributing significantly to the Roadrunner project, through performance analysis and modeling of the applications and system. The applications under considerations are, among others, Sweep3D, VPIC, Milagro, and Spasm. This work will contribute to the performance improvement of this system under a realistic workload and will be a major component of the upcoming assessment activities. The Performance and Architecture Lab's modeling methodologies will be employed to predict the performance of these applications on the full scale Roadrunner machine.
- Based on methodology developed by the Performance and Architecture Lab, significant improvements in system performance have been achieved through an application-based routing scheme for Infiniband. This will have a direct impact on the performance of Roadrunner at scale.
- Continued CellFS development for Roadrunner, which could increase programming productivity on the cell.
- The PaScalBB multi-lane architecture was demonstrated and put into production in the yellow network with two lanes and Redtail was brought up for acceptance, subsequently Redtail in the red network is now running with six lanes, with a growth path to 600 GB/sec. of input/output (I/O) bandwidth.
- Working closely with our global parallel file system (GPFS) vendor Panasas, we have the first production release of the Panasas file system, PanFS, that allows for the user the choice of redundancy schemes on a per file or directory basis. Previous testing has shown over a 15 times improvement in RAGE.
- Panasas file system has been successfully deployed on Redtail with greater than 3 GB/sec. provided to 1-connected unit (CU) job and 12 GB/sec. provided to a 4-CU job.
- Developed distance visualization techniques that progressively stream prioritized visual results to ASC scientists. These techniques immediately display and continuously update visual results; a significant improvement over standard visualization techniques that require waiting until the entire visualization process is completed before display.
- Developed comparative visualization software for analyzing and debugging ASC simulations by comparison and quantification of differences between ensembles of data.

SNL Accomplishments for Computational Systems and Software Environment

- Upgraded Red Storm. The original 10,368 compute nodes on Red Storm were upgraded to SeaStar 2.1 interconnect chips and 2.4-GHz Dual Core Opteron processors. A fifth row with SeaStar 2.1 interconnect chips and 2.4-GHz Dual Core

Opteron processors was integrated to increase the total number of compute nodes from 10,368 to 12,960. The upgraded Red Storm achieved 101.4 teraFLOPS on the Linpack benchmark and ranked second on the November 2006 TOP500 list.

- Developed the Catamount Virtual Node and deployed it into the Production Cray Code base to provide a seamless way for MPI (Message Passing Interface) applications to take advantage of the second core in the dual-core Opteron processors. Scaling studies on a wide range of applications showed excellent utilization of the second core.
- Released a scalable analysis tool (ParaView 3.0), with distributed Python scripting, a multiple view visualization capability, and a streamlined, customizable user interface.
- Developed a new environment, called the Immersive Topology Engine for Meshing, for guiding analysts through the process of generating a hexahedral or tetrahedral mesh for simulation. The Immersive Topology Engine for Meshing is built on the existing CUBIT Geometry and Meshing Toolkit and promises to dramatically improve the productivity of analysts in developing meshes for simulation. New geometric reasoning algorithms in the Immersive Topology Engine for Meshing detect potential problems in a computer aided design (CAD) model and provide a list of suggested solutions to the user in a wizard-like environment.
- Received Sigma 15 certification on the SNL simulation data management tools. These tools enable rigorous management of weapon analysis information and can now be used to manage sensitive use-control information.
- Demonstrated a prototype system to build and manage large ensembles of simulations that support UQ studies utilizing the SNL DAKOTA iterative systems analysis tool.
- Partnered the simulation data management team with the electrical system design tool developers to produce a full-system electrical repository for the rigorous management of weapon electrical models.
- Developed a lightweight and extensible job submission utility that homogenizes job submission processes across analysis codes and compute platforms by leveraging open source technologies.
- Provided support to the Level 1 Purple 100-teraFLOP milestone.

Facility Operations and User Support

LLNL Accomplishments for Facility Operations and User Support

- A major upgrade to the IBM Purple disk system was completed and the Purple system declared GA for completion of the Level 1 milestone. Procedures for operation of Purple as a tri-lab Computational Computing Campaign (CCC) resource were initiated. Utilization of Purple has regularly been about 90 percent.
- Completing a Level 2 milestone, the LLNL DCE infrastructure and services were retired from production use for Livermore Computing account and group management on both the unclassified and classified environments after over 10 years of operational use. The DCE/DFS file system was successfully decommissioned.
- ASCI White was decommissioned July 27, 2006, after more than five years as a production workhorse for the tri-lab ASC Program and SSP efforts.

- Site preparation, deployment, and GA for three new Linux InfiniBand-interconnected capacity clusters from the Peloton procurement (Rhea, Minos, and Hopi) were completed.
- Site preparation and delivery was completed for an expansion of the LLNL BlueGene/L system.
- Documentation and training materials were completed for the new LLNL capacity clusters and the new Moab resource manager. This activity was a deliverable for the Tripod Level 2 milestone.

LANL Accomplishments for Facility Operations and User Support

- Brought the Lightning system into full production with over 70 percent utilization. Lightning was the follow-on capacity system to assume work from the Q computers.
- Decommissioned half of the Q computing system, known as QA, on March 31, 2007.
- Delivered and accepted the Roadrunner base, which has a capacity of 70 teraFLOPS. The system was approved and accredited for secure operations in May, 2007, and selected weapons codes were run on the system in June 2007.
- Installed Tripod modules on several systems at LANL, in particular on the Roadrunner base system, as part of the tri-lab Level 2 milestone.
- Continued High Performance Computing (HPC) Division computer operator training for hardware maintenance of the ASC systems. Previously, hardware maintenance was handled by system engineers from the computer manufacturers. Operators now receive certification training from the vendor to assume this responsibility without voiding the manufacturer's warranties. Hardware support of additional clusters has been increased over the year to include Flash, Bolt, and Roadrunner. The team is currently responsible for servicing over 10,000 nodes.
- Installed and tested Ethernet network for the Roadrunner base system. Infiniband network management tools were developed to support the interconnect.
- Doubled the bandwidth for the DISCOM connection by installing two additional GigE TacLane encryptors. This increased the bandwidth to 4 GB/sec. The SecureNet encryptor was upgraded to a GigE encryptor. Worked with LLNL and SNL to improve file transfer performance across DISCOM. Increased file transfer rates for SNL to LANL transfers to over 200 MB/sec.
- Upgraded HPSS to v6.2 in both the red and yellow networks. The integration and deployment of STK/SUN T10000 tape and media technology occurred in both the red and yellow HPSS archive systems. Parallel storage interface (PSI) was successfully upgraded and deployed with the HTAR capability for bundling small files prior to transferring to the archive.

SNL Accomplishments for Facility Operations and User Support

- Achieved a 38-percent performance improvement for Thunderbird through software improvements to the Infiniband drivers and the MPI library, thus receiving a sixth rank on the TOP500 list with 53 teraFLOPS.
- Achieved GA status on Red Storm. New customers began Red Storm driver's license training, and Academic Strategic Alliance Program partners were issued accounts.
- Exceeded 30 terabytes of simulation files generated by Red Storm on data transfers from SNL to LLNL and LANL. Transfer rates to LANL achieved 400 megabytes/sec.

- Upgraded production HPSS systems to Version 6.2 and installed first of two new StorageTEK SL8500 tape libraries.

IV. Product Descriptions by the National Work Breakdown Structure

WBS 1.5.4: Computational Systems and Software Environment

The mission of this national sub-program is to build integrated, balanced, and scalable computational capabilities to meet the predictive simulation requirements of NNSA. It strives to provide users of ASC computing resources a stable and seamless computing environment for all ASC-deployed platforms, which include capability, capacity, and advanced systems. Along with these powerful systems that ASC will maintain and continue to field, the supporting software infrastructure that CSSE is responsible for deploying on these platforms includes many critical components, from system software and tools, to I/O, storage and networking, to pre- and post-processing visualization and data analysis tools. Achieving this deployment objective requires sustained investment in applied research and development activities to create technologies that address ASC's unique mission-driven need for scalability, parallelism, performance, and reliability.

WBS 1.5.4.1: Capability Systems

This level 4 product provides capability production platforms and integrated planning for the overall system architecture commensurate with projected user workloads. The scope of this product includes strategic planning, research, development, procurement, hardware maintenance, testing, integration and deployment, and quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include strategic planning, performance modeling, benchmarking, and procurement and integration coordination. This product also provides market research for future systems.

Capability Systems Deliverables for FY08

- Purple production capability cycles as per CCC guidance
- Sequoia Conceptual Baseline (CD-1) and Sequoia Performance Baseline/Construction Readiness (CD-2/3) Packages
- Assessment of Roadrunner Phase 3 final system and decision to move forward with Phase 3 procurement
- Given decision to proceed above, delivery and initiation of deployment of Roadrunner Phase 3 final system by Q4 FY08
- Provide maintenance for Red Storm to ensure it remains a useful and stable platform for ASC applications

WBS 1.5.4.1-LLNL-001 Purple

The Purple project has now delivered the IBM Purple cluster of pSeries POWER5 based SMPs at LLNL. The Purple system was declared GA in December 2006 and is currently

fully subscribed delivering CCC capabilities to the ASC customers. The expected system lifetime is five years. This project is currently in operational mode and receives funding for ongoing IBM maintenance.

Planned activities in FY08:

- Purple system maintenance as required

Expected deliverables in FY08:

- Continue to deliver production capability as per CCC guidance

Preliminary planned activities in FY09:

- Purple system maintenance as required

WBS 1.5.4.1-LLNL-002 Sequoia

The Sequoia project will acquire new computing resources focused on supporting UQ and reduction in phenomenology (the elimination of code "knobs"). The final Sequoia system will provide at least 12 times the sustained performance of ASC Purple on integrated design code calculations and 20 times the performance of BlueGene/L on a weapons science materials calculation. While Sequoia is to be sited at LLNL, final operation will fall under a national user facility paradigm similar to the Purple CCC and will be made available to LLNL, LANL, and SNL. In addition to a 2011 final system delivery, this project plans to provide a smaller, but significant, initial delivery (ID) environment (called Dawn), to be delivered in late calendar 2008 for needed scaling and code development to ensure effective use of the final platform. The Dawn ID environment will also be used for demanding near-term programmatic deliverables, in particular to support UQ and knob elimination.

There are two phases for Sequoia. Phase 1 (acquisition and delivery of Dawn ID environment) is to be completed by end of calendar 2008, and Phase 2 (acquisition and delivery of final Sequoia environment) is to be completed by end of calendar 2011. These environments (both Dawn ID and Sequoia final delivery) will consist of a large compute platform, plus requisite federated switch networking infrastructure and parallel file system storage hardware (similar to LLNL's existing Lustre parallel file system deployments) to support compute platforms. Acquired switching infrastructure and storage hardware may also have high-speed hardware connectivity to servers and resources at LLNL outside of Dawn and Sequoia, including visualization engines, archival storage movers, BlueGene/L, Purple, and TLCC07 Linux clusters.

Planned activities in FY08:

- Release of Sequoia request for proposal, evaluation of responses, and vendor selection
- Sequoia vendor contract negotiations and award approvals
- Planning for Dawn ID platform delivery

Expected deliverables in FY08:

- Sequoia Conceptual Baseline CD-1 package
- Sequoia Performance Baseline/Construction Readiness CD-2/3 package
- Sequoia final contract award

Preliminary planned activities in FY09:

- Sequoia ID installation and system integration

- Sequoia ID early science runs in the OCF
- Sequoia ID deployed in the secure computing facility (SCF) for SSP work

WBS 1.5.4.1-LANL-001 Systems Requirements and Planning

This project covers all aspects of program and procurement planning for future capability, capacity, and advanced systems and strategic planning for supporting infrastructure. The main focus of the project is to define requirements and potential system architectures for future capability platforms that meet ASC programmatic requirements and drivers. This project provides a focus for the various planning efforts and provides project management support for those efforts. This project will focus in FY08 on the project management and procurement of the hybrid Roadrunner phase 3 system to be deployed at LANL.

The current Roadrunner contract calls for a decision to be made in October 2007 on whether to exercise the option to procure the Phase 3 hybrid system from IBM. If the decision is made to proceed, then the focus of this effort will be to provide all the project management needed to deliver and deploy the Phase 3 system at LANL in FY08. The Phase 3 system is scheduled to deliver a significantly advanced architecture system that should provide compute power of over a petaFLOPS of computing cycles to the weapons program. The advanced architecture hardware will consist of a hybrid computing architecture that has the potential for significant improvements to the price/performance curve to help meet the computing requirements in the future.

The next phase for Roadrunner consists of deliveries for the assessment of the accelerator technology path leading to a decision to exercise the third phase procurement of the full set of advanced architecture technology. The third phase is the option for the acquisition of the final Roadrunner system. This system will consist of more than 3000 nodes interconnected via Infiniband 4x DDR. Each node will consist of dual core AMD Opterons connected to Cell Blades. If this option is exercised the delivery will be completed in FY08.

Planned activities in FY08:

- Pending an independent assessment, provide the overall project planning for the delivery and deployment of the Roadrunner phase 3 system. This will include system integration and management of the platform hybrid Phase 3 system. Since this project will provide overall planning and management of the Roadrunner project, actual execution will be by means of various projects in this Implementation Plan, as coordinated through the LANL Roadrunner project plan.
- Continue to prototype the application of the project execution model process for the procurement and integration of Roadrunner system. We are extending the process to have a full project execution plan for the entire project, which will be formally reviewed.
- Continue planning with the other labs to address perceived issues with the petascale computing environment. Personnel from our technical staff will continue to work with other tri-lab experts to address the petascale computing environment that will exist once petascale computational systems are available for tri-lab use in ASC.

Expected deliverables in FY08:

- Roadrunner Base system in full production at LANL
- Roadrunner Phase 3 final system delivered and initial deployment started

Preliminary planned activities in FY09:

- Deployment of Roadrunner Phase 3 into the network environment

WBS 1.5.4.1-SNL-001 Red Storm Capability Computing Platform

Red Storm is a tightly coupled massively parallel processor compute platform with approximately 125 teraFLOPS of peak processing capability. The machine uses 2.4 GHz dual-core AMD Opteron processors and a custom, very high performance, 3D-mesh communication network. Red Storm has a total of 13,600 dual-core Opteron processors and over 35 terabytes of memory and 400 terabytes of high-performance local disk that is split equally between classified and unclassified use. Red Storm produced a 101.4 teraFLOPS on the HPL benchmark and took over the number two position on the November 2006 Top 500 list of the world's fastest supercomputers. Cray now has over 18 sites and has 35 systems sold based on the Red Storm architecture.

Specifications and an expanded description are available at http://www.sandia.gov/ASC/redstorm.

Planned activities in FY08:

- During the first quarter of FY08, SNL will integrate over 1.0 petabytes of additional disk storage into the Red Storm system for the classified end. SNL jointly with Cray will continue to improve both the reliability of and performance of the Red Storm system throughout FY08.

Preliminary planned activities in FY09:

- Red Storm will be in full production throughout FY09

WBS 1.5.4.2: Capacity Systems

This level 4 product provides capacity production platforms commensurate with projected user workloads. The scope of this product includes planning, research, development, procurement, hardware maintenance, testing, integration and deployment, and quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include the procurement and installation of capacity platforms.

Capacity Systems Deliverables for FY08

- Full acceptance and production deployment at LLNL, LANL, and SNL of initial TLCC07 SUs
- Full production deployment of capacity Roadrunner Base system at LANL
- Deployment of the Tripod Operating System Software (TOSS) on TLCC systems
- Support TLCC procurement process
- Support deployment of the TOSS stack on TLCC systems

WBS 1.5.4.2-LLNL-001 Tri-Lab Linux Capacity Cluster

The TLCC07 procurement will provide the ASC community with a large amount of capacity computing resources in FY08. To affordably and efficiently provide this production quality computing capacity, the TLCC07 project has embarked on an approach that extends existing tri-lab practices and leverages extensive experience fielding world class Linux clusters within the tri-lab community.

TLCC07 represents a new strategy for quickly building, fielding and integrating many Linux clusters of various sizes into classified and unclassified production service through a concept of SUs. The programmatic objective is to dramatically reduce the overall total cost of ownership of these "capacity" systems relative to the best practices in Linux cluster deployments today. This objective only makes sense in the context of these systems quickly becoming robust and useful production clusters under the crushing load that will be inflicted on them by ASC and SSP scientific simulation capacity workloads.

The TLCC07 procurement is designed to maximize the purchasing power of the ASC program by combining all tri-lab ASC capacity computing procurements into a single procurement. It is anticipated that purchasing a huge set of common hardware components will lead to lower cost through high volume purchases. In addition, by deploying a common hardware environment multiple times at all three sites, it is anticipated that time and cost associated with any one cluster will be greatly reduced. The common TLCC07 hardware environment will also accelerate the adoption of a common tri-lab software environment (Tripod).

Planned activities in FY08:

- SU build and testing phase to coordinate with the selected TLCC07 vendor at the vendor site. The SUs must pass all the workload tests in the contract criteria prior to moving on to the next phase/activity.
- SU delivery and acceptance phase to coordinate TLCC07 vendor at LLNL. This testing exactly replicates the testing and burn-in done at the vendor site to ensure no equipment was damage during cluster breakdown or shipping.
- SU cluster integration phase to collaborate with TLCC07 vendor to aggregate SUs into multi-SU clusters of 2, 4, 6, or 8 SUs. The multi-SU cluster will be integrated into the LLNL simulation environment and complete the benchmark and application acceptance tests.
- Complete the I/O and networking infrastructure procurements for TLCC07.
- Retire the ALC OCF cluster.
- Retire the Lilac SCF capacity cluster, which will eliminate support of x86 production compute architectures at LLNL.

Expected deliverables in FY08:

- Full acceptance and production deployment of initial TLCC07 SUs running TOSS

Preliminary planned activities in FY09:

- Upon the direction of NNSA ASC Headquarters, the TLCC technical team will explore a follow-on TLCC capacity procurement in FY09–FY10

WBS 1.5.4.2-LANL-001 Capacity Integration

The Capacity System Integration project will continue system integration for all capacity systems at LANL and will focus on continuing the integration of the Roadrunner Base System into full production, integrating the TLCC07 and TLCC08 systems delivered to LANL in FY08, and continue to support other capacity systems at LANL. The Roadrunner Base system is a 70-teraFLOP machine delivered to LANL in FY07 for capacity computing work and is also a precursor to the Roadrunner Phase 3 final system to be delivered in FY08. The TLCC is a capacity system that will be procured using the July 2007 request for proposal describing a system located at each of the three weapons

labs. The system integration effort will address complex integration issues relating to high-end terascale computing environments and provide the necessary resources to stand up such systems for production computing work.

Planned activities in FY08:

- Continued coordination of integration and testing of the Roadrunner Base System for limited access to application codes evolving to full production and general availability in early FY08.
- Plan for the deployment of TLCC systems at LANL. The first systems, TLCC07, are targeted for the open Turquoise network at LANL. Subsequent systems, TLCC08, will be in the secure partition.

Expected deliverables in FY08:

- Roadrunner Base System in full production and general availability
- Full acceptance and production deployment at LANL of initial TLCC07 SUs
- Integrate Tripod selected stack and tools deployed on TLCC cluster (with Application Readiness Team)

Preliminary planned activities in FY09:

- Upon the direction of NNSA ASC Headquarters, the TLCC technical team will explore follow-on TLCC capacity procurements in FY09–FY10

WBS 1.5.4.2-SNL-001 Deployment of ASC Capacity Systems

The purpose of this project is to support the acquisition of new ASC capacity systems that minimize the total cost of ownership for capacity computing performance. This project coordinates the integration and deployment of TLCC capacity systems into SNL's production computing environment in collaboration with WBS 1.5.5.1-SNL-001 Production Computing Services.

This project is supported by analysis of SNL's portfolio of application needs for capacity computing systems within the context of the higher integrated ASC platform strategy of capability, capacity, and advanced systems. Efforts include definition of requirements for TLCC system procurements and support for Tripod common software stack for new and existing capacity systems.

Planned activities in FY08:

- Continue to support tri-lab ASC TLCC07 system procurement process
- Identify the appropriate integration of Tripod system software with the broader open source software community tools and capabilities currently deployed on SNL's commodity Linux cluster capacity systems

Expected deliverables in FY08:

- Support TLCC procurement process
- Deploy the ASC Tripod common software stack on TLCC systems

Preliminary planned activities in FY09 include:

- Support preparation of the TLCC10 statement of work and competitive request for proposal for the next generation of ASC Capacity Computing systems
- Retiring obsolete SNL Linux clusters by replacing them with TLCC clusters

WBS 1.5.4.3: Advanced Systems

This level 4 product provides advanced architectures in response to programmatic, computing needs. The scope of this product includes strategic planning, research, development, procurement, testing, integration and deployment, as well as industrial and academic collaborations. Projects and technologies include strategic planning, performance modeling, benchmarking, and procurement and integration coordination. This product also provides market research, and the investigation of advanced architectural concepts and hardware (including node interconnects and machine area networks) via prototype development, deployment and test bed activities. Also included in this product are cost-effective computers designed to achieve extreme speeds in addressing specific, stockpile-relevant issues through development of enhanced performance codes especially suited to run on the systems.

Advanced Systems Deliverables for FY08

- BlueGene/P demonstration system to be delivered at Argonne National Laboratory
- Specific recommendations for how to improve the balance factors that will limit scalability of systems available in the 2012–2016 timeframe
- Assuming that the decision has been made to proceed in October 2007, deliver the Roadrunner Phase 3 system with the goal to achieve petaFLOP-level computing performance to meet ASC weapons computing requirements
- A report detailing work completed in support of a high-speed interconnect subsystem and an advanced memory subsystem to address the needs of ASC applications in the 2012–2020 timeframe

WBS 1.5.4.3-LLNL-001 BlueGene/P and BlueGene/Q Research and Development

This project is a multi-year NNSA and Office of Science R&D partnership with IBM on advanced systems that target the development and demonstration of hardware and software technology for 1-petaFLOPS and 10-petaFLOPS systems. The BlueGene/P hardware is based on an extension of the highly successful BlueGene/L architecture with faster nodes, more memory, faster interconnects and larger system scalability. The software approach to BlueGene/P is open-source collaborative development between IBM research, Linux Technology Center, E&TS division and Argonne National Laboratory and the ASC tri-labs. In FY08, a BlueGene/P demonstration system is expected to be delivered to Argonne. The follow-on BlueGene/Q system design targets a 10-petaFLOPS system at the end of the contract.

This project incorporates requirements from the DOE laboratories, especially Argonne and LLNL, to have input into the design choices and system testing for microprocessors, node architectures, and interconnects.

Planned activities in FY08:

- BlueGene/Q reviews, including core network design, system packaging, ASICs, and overall architecture

Expected deliverables in FY08:

- BlueGene/P demonstration system delivered at Argonne

Preliminary planned activities in FY09:

- BlueGene/Q hardware component releases for ASIC design and system packaging

- BlueGene/Q compiler review

WBS 1.5.4.3-LLNL-002 Petascale Application Enablement

This project enables advanced application work to develop benchmarks for new platforms, such as Sequoia, and to adapt current codes to the expected new architectures. A primary target of this project is investigating ways to improve application thread performance for future many-core platforms. The project team efforts include both direct application work and benchmark development and testing.

Planned activities in FY08:

- Study of application performance requirements for Sequoia
- Investigation of opportunities for thread-parallel performance in production applications

Preliminary planned activities in FY09:

- Science and weapon code testing on Sequoia ID (Dawn) system, especially measuring single node thread performance
- Continued focus on code improvement opportunities

WBS 1.5.4.3-LANL-001 Roadrunner Phase 3 Procurement

The Roadrunner Phase 3 Procurement is an option in the Roadrunner contract. A final system assessment is scheduled for October 2007. Based on this technical assessment, a decision will be made on the procurement of the Phase 3 system. The Roadrunner base system was delivered, installed, accepted, and has passed an ASC Level 2 milestone for applications use of the system (June 2007). The final system is configured with hybrid nodes based on a hybrid architecture using IBM System AMD Opteron-based processors accelerated with IBM's Cell Broadband Engine blades.

Hybrid computing architectures are an important direction for high-performance computing. This project will look at the software and programming issues relative to hybrid computing by looking at algorithms and applications and their applicability and performance enhancements on this advanced architecture system.

The projected peak performance of the final system is approximately 1.4 petaFLOPS.

Planned activities in FY08:

- A final system technical assessment will be conducted in October 2007. This is part of a Level 2 milestone for FY08.
- Procure a hybrid advanced architecture for the final Roadrunner system. Several software deliveries are also scheduled that will provide the full system with a hybrid architecture using the IBM System AMD Opteron-based host servers accelerated with IBM's Cell Broadband Engine blades. This is part of a Level 2 milestone in FY08.
- Full scale demonstration run at the vendor site.
- Delivery of system to LANL.
- Acceptance testing of Roadrunner Phase 3 system.

Expected deliverables in FY08:

- Roadrunner Phase 3 final system built and tested at the vendor site and then delivered to LANL for integration and acceptance testing

Preliminary planned activities in FY09:

- The system will be integrated into the LANL computing environment for initial testing using selected science applications; the system will then be transitioned to the LANL computing environment

WBS 1.5.4.3-SNL-001 Advanced Systems Technology Research and Development

This project is focused on the research and development of specific capabilities to actively address the anticipated imbalances between processors and their associated memory and high-speed interconnect subsystems in the 2012 to 2020 timeframe. SNL will use this project to lead the development of an interconnect infrastructure to address the anticipated needs of future generations of capability supercomputers. The interconnect work will develop requirements and specifications for both the network interface controller and associated router subsystem to satisfy the desired levels of network performance.

In addition, this project is investigating alternative memory subsystem architectures to address processing bottlenecks in future scientific computing processor architectures. The goal of this project is to be proactive in working with the commercial processor and memory vendors to develop new memory architectures and standards to ensure efficient utilization of the computational resources in future generations of commodity processors.

This project is also investigating planned future capability class systems and advanced architectures and their impact on the performance of SNL's application code base. By understanding the impact such architectures will have on application performance, we will be able to make informed choices on future platform acquisitions.

Planned activities in FY08:

- R&D requirements and desired functionality for future high-speed interconnect subsystems
- R&D requirements and desired functionality for future processor memory subsystems
- Evaluation of the performance of SNL applications on future capability and advanced system architectures

Expected deliverables in FY08 include:

- A report detailing work completed in support of a high-speed interconnect subsystem and an advanced memory subsystem to address the needs of ASC applications in the 2012 to 2020 timeframe

Preliminary planned activities in FY09:

- Continued work in high-speed interconnect and memory subsystems to address future computational needs
- Work towards satisfying the "Advanced Memory Subsystems" Level 2 milestone

- Work towards satisfying the "Evaluation of the Impact Chip Multiprocessors have on SNL Application Performance" Level 2 milestone

WBS 1.5.4.4: System Software and Tools

This level 4 product provides the system software infrastructure, including the supporting operating system environments and the integrated tools to enable the development, optimization and efficient execution of application codes. The scope of this product includes planning, research, development, integration and initial deployment, continuing product support, and quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include system-level software addressing optimal delivery of system resources to end-users, such as schedulers, custom device drivers, resource allocation, optimized kernels, system management tools, compilers, debuggers, performance tuning tools, run-time libraries, math libraries, component frameworks, other emerging programming paradigms of importance to scientific code development and application performance analysis.

System Software and Tools Deliverables for FY08

- Deploy system software and tools environment for TLCC07 and Sequoia ID
- Deploy MOAB resource manager on BlueGene/L
- Certify TLCC07 Tripod applications development tools environment
- Deploy initial highly scalable code correctness tool suite
- Contribute to Level 2 milestone "ASC Petascale Environment Infrastructure Deployment Plan" planning document
- Detailed assessment and plan for petascale environment requirements
- Deploy Tripod stack and tools on TLCC system at LANL
- Deploy MOAB on Roadrunner and TLCC systems at LANL
- Report on performance modeling and measurement verification to support the Roadrunner Phase 3 assessment and build up and integration of the full Phase 3 system
- Report on hybrid programming successes and issues to the Roadrunner Phase 3 assessment
- Provide OpenMPI support for Cell Blade to Opteron communications and job launch over PCIe link
- Provide OpenMPI enhancements for performance and scalability on the Roadrunner Base and Cell-accelerated systems
- Work with IBM to provide a hybrid debugger capability across multiple Cell-accelerated nodes
- Define a set of performance metrics for each code project
- Enhance performance of X Division codes at each major release
- Design and prototype initial Domain Specific Language for hydro; work with X Division to specify Domain Specific Language
- Complete the integration of LANL's Fortran parser with the ROSE compiler infrastructure

- Prototype cross-domain infrastructure for domain-specific languages and libraries
- Report on implementation of a scalable streaming computation, which uses network processors
- Prototype a set of program correctness tools which aid programmers in creating and maintaining correctly performing codes
- Use application kernel performance models to examine the expected performance when scaling the number of advanced devices
- Develop a performance testing suite for the identification of common and idiosyncratic performance characteristics
- Provide system and runtime software to address the technical issues of paramount concern to running ASC applications on our capability systems with sufficient scalability, reliability, and application performance to meet our SSP requirements
- Tools to measure reliability, availability, and serviceability (RAS) performance of ASC systems
- Updated benchmark measurements of top priority applications on the existing portfolio of ASC platforms
- Deploy the O|SS performance tool on Red Storm
- Develop tools that process application performance data (such as MPI traces and hardware performance counter data) into a form that can be used to perform computer architecture simulations
- Develop and implement processes required for continued collaborative development and deployment of TOSS

WBS 1.5.4.4-LLNL-001 System Software Environment for Scalable Systems

This project provides system software components for all the major platforms at LLNL, research and planning for new systems and future environments, and collaborations with external sources such as the platform partners, especially IBM and Linux vendors. This project covers system software components needed to augment Linux and required proprietary operating systems that function in a manageable, secure, and scalable fashion needed for LLNL platforms. Currently, this includes work on developing, modifying, and packaging the TOSS, and developing scalable system management tools to support the operating system and interconnect (for example, TOSS and InfiniBand monitoring tools), as well as the resource management environment to queue and schedule code runs across LLNL systems.

Planned activities in FY08:

- Initial definition, deployment, and support of CHAOS-based TOSS at all three labs
- Ongoing SLURM, TOSS, LCRM, and MOAB software development and support
- Preparation for system software environment for Sequoia ID and TLCC07 SUs
- Preparation for an instantiation of the Tripod RAS metrics implementation
- Enhancements to host monitoring software to improve ticketing of system events
- Develop InfiniBand performance manager software and deploy on TLCC07 systems
- Work with Cluster Resources to integrate Gazebo into TOSS via MOAB

- Work with Cluster Resources to improve security of MOAB

Expected deliverables in FY08:

- Deploy TOSS, SLURM and MOAB on TLCC07
- Deploy TOSS, derived from RHEL 5, on all LLNL Linux clusters including TLCC07
- Deploy MOAB on BlueGene/L

Preliminary planned activities in FY09:

- Deploy system software environment on Sequoia ID system
- Ongoing SLURM, TOSS, LCRM, and MOAB software development and support
- Improve the Tripod system operating environment

WBS 1.5.4.4-LLNL-002 Applications Development Environment and Performance Team

The Applications Development Environment and Performance Team project provides the code development environment for all major LLNL platforms, supports user and code productivity, provides research and planning for new tools and future systems, and collaborates with external sources of code development tools such as platform partners, independent software vendors, and the open source community. The project works directly with code developers to apply tools to understand and to improve code performance and correctness. The project resolves bug and user trouble reports, including interactions with the software providers to fix problems.

The elements of the development environment covered by this project include, but are not limited to, compilers, debuggers, performance assessment tools and interfaces, memory tools, interfaces to the parallel environment, code analysis tools, and associated run time library work, with explicit focus on the development environment for large-scale parallel platforms. It does not cover the development environment for workstations or other special purpose systems.

The project has established expertise in research, development and application of code development environments for large-scale systems. It has developed or extended several components of the installed production environment. Integration of project members with code development teams ensures high performance use of existing systems and supports customer-based planning of future improvements to the environment. Similarly, long-term relationships with external partners, such as IBM, TotalView Technologies and Openworks, ensure that project members can resolve trouble reports quickly and avoid unnecessary duplication of existing capabilities.

Planned activities in FY08:

- Provide and maintain Purple and BlueGene/L code development environments
- Coordinate integrated design code scaling for Sequoia
- Provide and refine the common tri-lab environment for TLCC07
- Identify, assess, and specify petascale development environment approaches
- Develop new techniques to improve robustness and performance of ASC codes
- Interact with the ASC code teams and vendors to improve software products

Expected deliverables in FY08:

- Detailed assessment and plan for petascale environment requirements
- Certification of the TLCC07 Tripod Applications Development Tools environment
- Deployment of initial release of highly scalable code correctness tool suite

Preliminary planned activities in FY09:

- Continue code development environment support on all LLNL ASC platforms
- Identify and develop refinements of the code development environment for existing and future capacity and petascale systems
- Continue to support users and interact with vendors to serve user needs

WBS 1.5.4.4-LANL-001 Roadrunner Computer Science

This project supports the success of the full Roadrunner system, from the integration and performance of the base system through the support of the hybrid system running at a sustained petaFLOPS. Perform work in area of systems, Open MPI, performance measurement, analysis and modeling, tools, and architecture for the Roadrunner project during the planning, initial hardware availability, build-up, and pre- and post-installation of the system.

Planned activities in FY08:

- Develop OpenMPI support over PCIe links between Opteron hosts and Cell blades of Roadrunner Phase 3 system
- Develop full application, algorithm, and Linpack performance models
- Support IBM Linpack performance modeling efforts
- Performance tuning of Roadrunner base system and the Roadrunner hybrid system in production

Expected deliverables in FY08:

- Reports on performance modeling and measurement verification to support the Roadrunner Phase 3 assessment and build up and integration of the full Phase 3 system
- Reports on hybrid programming successes and issues to the Roadrunner Phase 3 assessment
- Provide OpenMPI support for Cell Blade to Opteron communications and job launch over PCIe link
- Work with IBM to provide a hybrid debugger capability across multiple cell-accelerated nodes

Preliminary planned activities in FY09:

- This element will transition into the Usable Supercomputer Project

WBS 1.5.4.4-LANL-002 Usable Supercomputer Project

The scope of this project is to measurably improve the usability, performance, reliability, efficiency, and productivity of petaFLOPS supercomputers, beyond the Roadrunner project, for nuclear weapons applications. The main target for this fiscal year will be multicore and hybrid designs, beyond the Roadrunner project, which will include

WBS 1.5.4.4-LANL-003 Code Performance and Throughput

The goals of this project are to measurably improve the performance of the integrated X Division weapons code projects at each major release, provide predictive performance tools for the same projects, and inform the code re-factorization project. The Code Performance and Throughput project will develop performance prediction tools for the tri-lab user community while meeting specific requirements for the hybrid Roadrunner system. To ensure relevance, this project will deploy staff directly on code projects to enhance code performance in each major release. As an overarching goal the project will seek to estimate the runtime based on large-scale variables like physics method, computer system, data sets, and other factors. This project will define targets for enhanced performance as a function of time, deploy computer science expertise to both major X Division weapons codes projects, and enhance the performance of the codes against those targets as delivered to the user community.

Planned activities in FY08:

- Define a set of performance metrics to measure for each code project
- Place measurement of performance using these metrics into nightly code regression testing
- Profile codes and look for hot spots for improvement
- Define a set of performance targets to hit at each major release during the fiscal year (for example, 30-percent improvement in performance at first release, 10 percent at second, and so on)
- Design and develop a performance prediction tool for the user community (this physics + this computer + … = this time to run problem)
- Provide input to the re-factorization project

Expected deliverables in FY08:

- Define a set of performance metrics for each code project
- Enhance performance of X Division codes at each major release

Preliminary planned activities in FY09:

- This activity will extend their purview to additional X division weapons code activities in FY09

WBS 1.5.4.4-LANL-004 Tripod Software Development/Deployment

The Tripod Steering Committee, representing all tri-lab sites, is identifying tools and a software stack for deployment on the tri-lab cluster computers that are to be delivered in FY08. The goal is to control costs of software development by having a common framework, tools, and software stack on all TLCC systems at all tri-lab sites.

Planned activities in FY08:

- Develop and deploy Tripod software modules on the LANL TLCC cluster. The tri-lab compatible software stack components are to include: operating system, communication libraries, compilers, application development tools, system management tools, and user environment. The tri-labs will continue to do gap and risk analyses of the Tripod software stack and develop a mitigation plan accordingly. Configuration management and change control processes will be developed as tools to manage the Tripod stack to remain consistent across the tri-labs.

Expected deliverables in FY08:

- Tripod stack and tools deployed on TLCC systems at LANL.
- Design and implement a baseline application monitoring tool able to accurately determine if applications are making progress. Test and deploy a beta-version of the baseline implementation on TLCC.
- Open|SpeedShop: Integration of a Javelina plug-in into OSS (code coverage); work with SNL on hardening build environment and user support.
- Lead analysis and deployment of the sharing environment. If Sourceforge at LANL is selected then establish connection, authentication and use processes.
- Gazebo: Enhanced charting tools for test results analysis.
- Gazebo: Continued enhancement and maintenance of the existing Gazebo suite.
- Gazebo: Incorporation of the CBENCH software-testing suite from SNL as a modular component of building and deploying tests.
- Gazebo: Refinement of the acceptance testing component of Gazebo as the TLCC acceptance test package, tuned according to the first experiences with the TLCC delivery.

Preliminary planned activities in FY09

- Test and deploy a release-version of the baseline implementation on TLCC and other platforms as appropriate. Design and implement advanced application monitoring functionality integrated into the baseline tool. Test and deploy beta-version of the advanced implementation as appropriate.

WBS 1.5.4.4-SNL-001 Software and Tools for Scalability and Reliability Performance

This project supports software development to address scalability and reliability performance on SNL computational systems. This project has three components: 1) Red Storm system software development and support, 2) development of RAS performance capabilities, and 3) application performance benchmarking and analysis for current and future ASC systems.

The Red Storm system software development effort provides the core support for Red Storm's Catamount Virtual Node lightweight kernel and runtime software. New efforts are underway to improve support for Python-enabled ASC applications. This project also supports efforts to identify and resolve Red Storm I/O and parallel file system reliability issues and performance bottlenecks.

The RAS performance capabilities are focused on two areas: First, the development of standardized RAS metrics and tools to instrument current ASC capability and capacity systems (requires collaboration with and handoff to FOUS WBS 1.5.5.1-SNL-001 Production Computing Services). Second, RAS subsystem design specification to prepare for future ASC system designs. This activity will also support the development of formal specifications and design for a RAS subsystem to help ensure future capability systems have better RAS functionality and performance than Red Storm.

The application performance benchmarking, analysis, and modeling provides the ability to assess component and sub-system tradeoff for future system architectures, and guide the mapping of our workload of computational work to the ASC portfolio of platform resources.

emergent designs (such as the LLNL Sequoia platform). Given that system design and performance cannot be decoupled from the application workload that runs on such a system, the project will examine the multiple challenges of extreme-scales for workloads of interest to the weapons program. Experimentation, measurement, quantitative analysis and modeling, and software development will be employed for defining a supercomputing environment optimal for the targeted application workload, leading to high productivity in all of its facets.

The work in this project will provide direct guidance to code developers and system vendors, including the development of hardware and software prototypes for the future. The work will be tightly integrated (and is related to) the Roadrunner computer science project, the Roadrunner weapons science project, the Roadrunner algorithms project, the code performance and throughput project, the accuracy project, and the integrated codes project. Significant interaction is anticipated with multiple HPC projects in CSSE and FOUS.

Planned activities in FY08:

- New programming methodologies for heterogeneous multi-core architectures
- Evolutionary code acceleration techniques for commodity processors
- Tools and techniques for increasing fault-tolerance and robustness of numerical codes
- Scalable streaming computation using network processors
- Static analysis and source-to-source transformation tools to run across multiple hardware computer architectures
- Virtualization techniques for heterogeneous architectures
- Dynamic performance models for application steering for performance

Expected deliverables in FY08:

- Design and prototype initial Domain Specific Language for hydro work with X-Division to specify Domain Specific Language
- Implement accelerated kernels and report on demonstrated performance improvements
- Complete the integration of LANL's Fortran parser with the ROSE compiler infrastructure
- Implement fault tolerant compact app hydro code in fault tolerant runtime system
- Prototype cross-domain infrastructure for domain-specific languages and libraries
- Prototype a set of program correctness tools which aid programmers in creating and maintaining correctly performing codes
- Use application kernel performance models to examine the expected performance when scaling the number of advanced devices
- Develop a performance testing suite for the identification of common and idiosyncratic performance characteristics

Preliminary planned activities in FY09:

- Extend debugger and tools capabilities on Roadrunner
- Support for ongoing Science@Scale efforts and advanced architecture development and integration of prototype capabilities into integrated codes for early testing

Planned activities in FY08:

- **C**ontinued support for the upgrade of Red Storm to dual-core Opterons with Catamount Virtual Node.
- Establish standardized metrics for measuring RAS performance and develop tools that allow RAS performance measurements on all our systems. Through Tripod, the RAS Metrics project is working with counterparts at LANL and LLNL.
- Continued efforts to address system software performance and reliability issues; and benchmarking our application performance on the evolving portfolio of ASC platforms.

Expected deliverables in FY08 include:

- System and runtime software to address the technical issues of paramount concern for running ASC applications on our capability and capacity systems.
- Support an anticipated task to produce a RAS specification and metrics implementation as part of the FY08 Tripod Level 2 milestone. The SNL–specific deliverables for this product will include tools to measure our RAS performance.
- Guidance to the Capability Planning Advisory Committee and to the SNL Platform Oversight Committee on the best mapping of our application workload to our (evolving) portfolio of ASC systems based on our updated application benchmark measurements.

Preliminary planned activities in FY09:

- Extension of Catamount to Quad-core Opteron support.
- Expansion of standardized RAS metrics beyond ASC to the broader HPC community to allow common tools to be developed that analyze machine failures, locate marginal hardware, and inform job owners of possible machine causes for job crashes.

WBS 1.5.4.4-SNL-002 ASC Programming Tools Project

This project is concerned with providing programmers tools to aid them in making their code reliable and efficient. These tools will be designed to be useful for collecting application data for use in architecture design and tuning. Programming tools for the ASC program have unique requirements including coexistence with MPI, scalability to tens of thousands of processors, and operability on ASC platforms.

The tools required fall broadly into two categories, performance analysis tools and debugging tools. Performance analysis tools can be separated into tools for examining how well a program utilizes a processor and tools for examining how efficiently a program uses the communication network. Debugging tools can be divided into traditional debuggers, and memory debugging tools.

Planned activities in FY08:

- Participate in the oversight of tri-lab Open|Speedshop and Valgrind contracts for tool development
- Explore better ways to characterize application performance
- Selectively deploy preliminary tool releases to code teams for evaluation and testing
- Continually familiarize ourselves with existing and new tools, developing training materials and usage guidelines for the tools

FY08 expected deliverables include:

- Deployment of the O|SS performance tool on Red Storm
- Development of tools that process application performance data (such as MPI traces and hardware performance counter data) into a form that can be used to perform computer architecture simulations
- Oversight of the O|SS and Valgrind contracts
- Further development of enhanced performance diagnostics

Preliminary planned activities in FY09:

- Review successes and failures of the latest tool features based on users' experiences and planning work to correct problems, as well as continued contract oversight

WBS 1.5.4.5: Input/Output, Storage Systems, and Networking

This level 4 product provides I/O (input/output, or data transfer) storage infrastructure in balance with all platforms and consistent with integrated system architecture plans. The procurement of all supporting subsystems, and data transfer, storage systems and infrastructures occurs through this product. The scope of this product includes planning, research, development, procurement, hardware maintenance, integration and deployment, continuing product support, and quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include high-performance parallel file systems, hierarchical storage management systems, storage-area-networks, network-attached storage (NAS), and HPSS or future hierarchical storage management system disks, tape, robotics, servers, and media. This product also includes relevant prototype deployment and test bed activities. Projects and technologies in the advanced networking and interconnect areas shall include networking and interconnect architectures, emerging networking hardware technologies and communication protocols, network performance/security monitoring/analysis tools, and high performance encryption and security technologies.

Input/Output, Storage Systems, and Networking Deliverables for FY08

- Release HPSS R7.1 and operational plan for FY09 production deployment
- Use in production new SL8500 tape libraries
- Procure and deploy archival storage hardware and software to support expansion of ASC platforms
- Deploy of Lustre release 1.8
- Evaluate and analyze Zettabyte File System (ZFS) for prototype ZFS replacing ext3 underlying Lustre file system (or possible backend use on NAS servers)
- Contribute to Level 2 Milestone "ASC Petascale Environment Infrastructure Deployment Plan" planning document
- Deliver and help with deployment of HPSS v7.1
- Provide new version of PSI with enhanced HTAR support
- Collect requirements for next releases of HPSS and PSI
- Create plans for implementing user requirements
- Design, develop, and test next versions of HPSS and PSI

- Deploy dead I/O node detect and reroute on all 3 GPFS environments red, yellow, turquoise
- Demonstrate initial PaScalBB switch and storage level fail over (not deploy just demonstrate - will require Panasas alternate pathing from shelf) (full fail over end to end for I/O)
- Deploy I/O infrastructure to RR phase 3 hybrid system
- Provide scalable network file system (NFS) single name space service report to assist in decision for deployment
- Write I/O section for the Level 2 petascale infrastructure milestone
- Lead I/O research for HPC systems (POSIX standards committee, parallel NFS, Northwestern University MPI-IO work, and Lightweight File System)
- Test and optimize the MPI-IO, parallel HDF5, and NetCDF libraries on Red Storm and the RoSE clusters; testing support for end-to-end I/O optimization

WBS 1.5.4.5-LLNL-001 Archive Storage

The Archival Storage Project provides end-to-end long-term, high performance, archival storage services to ASC customers. This includes a collaborative software development effort between the tri-labs, ORNL and IBM as well as deployment, and support of archival storage software and interfaces for tri-lab ASC customers on unclassified and classified networks. It also includes the selection, procurement, deployment, support, and maintenance of archival storage hardware and storage media and the ongoing technology refresh and data stewardship.

Archival storage system software (currently, HPSS) provides scalable, parallel archival storage interfaces and services to customers running at all three NNSA/ASC labs. HPSS distributes data across a configurable amount of storage units and removes other limits to scaling including number of files, directories, and concurrent users.

A world-class array of storage hardware is integrated beneath HPSS, supplying the performance necessary to offload ASC platforms and not hinder computation. This includes disk arrays, tape subsystems, mover nodes, storage-area-networks, networks, robotics and petabytes of media. Together, this hardware and software supports unlimited storage for an unlimited amount of time, currently at speeds in excess of 2.5 GB/sec.

Planned activities in FY08:

- Complete the development of HPSS R7.1 through code review, system test, integration test and final release. HPSS R7.1 is focused on file create and transactional performance in support of petascale environments.
- Investigate requirements for deployment of HPSS R7.1 in production environments, identifying and procuring requisite infrastructure components for example, new 64-bit core server platform, metadata disk, and DB2 v9.
- Test and integrate new SL8500 enterprise-class tape libraries into production archives to enable retirement of aging tape technologies, in order to decommission aging robotic tape silos by end of CY2010, when support ends.
- Provide ongoing support of currently deployed archival storage systems, including selection, deployment, support and maintenance of all archival storage hardware and media, customer and interface support, and ongoing tech refresh, and data stewardship of LLNL archives.

Expected deliverables in FY08:

- Release of HPSS R7.1
- Operational plan for FY09 production deployment of HPSS R7.1
- Production use of new SL8500 tape libraries
- Archival storage hardware and software procured and deployed to support expansion of ASC platforms

Preliminary planned activities in FY09:

- Deploy HPSS R7.1 in production archives, improving small file create performance and transaction rates for petascale environments in direct support of a planned Level 2 milestone for FY09
- Launch the tri-lab development process for HPSS R8.1, focusing on dynamically scaling core servers for a petascale environment

WBS 1.5.4.5-LLNL-002 Parallel and Network File Systems

This project provides for the development, testing (feature, capability, performance, and acceptance), procurement, integration, and ongoing support of various file system technologies and interfaces necessary for the efficient and effective use of ASC high-performance platforms. Included is the continuing development and support of Lustre as a fully-featured file system for the range of ASC capability and capacity platforms, the deployment and support of GPFS on the ASC IBM platforms, the deployment and support of ubiquitous NAS services for home, project, and scratch space, and the I/O support of various programming interfaces for parallel I/O.

This project deploys and supports Lustre and GPFS file systems for ASC platforms as well as high-availability NAS file systems for home and project space, and scratch space for serial capacity clusters. The file system up through the programming interfaces are supported to help developers of applications use parallel I/O effectively.

Planned activities in FY08:

- Provide Lustre file system performance and scalability enhancements. Further, investigate and provide any necessary file system enhancements related to Sequoia ID, Sequoia, and TLCC platforms.
- Begin deploying NFSv4 as a replacement for current NFSv3-based NAS services while continuing ongoing support of currently deployed NAS services and investigating the expansion of NAS data redundancy and backup technologies.
- Analyze and evaluate Sun's ZFS prototyped for Linux as a replacement for ext3 as the underlying file system for Lustre as well as for backend use on NAS servers for scratch space.
- In support of tri-lab FY08 Level 2 petascale environment milestone, specify I/O, file system, and storage components necessary for petascale environment and provide formal planning document.
- Maintain GPFS and Lustre parallel file system support, and middleware and higher-level I/O library support for users.

Expected deliverables in FY08:

- Deployment of Lustre release 1.8 across OCF and SCF Lustre file systems

- ZFS evaluation and analysis for prototyped ZFS replacing ext3 underlying Lustre file system or possible backend use on NAS servers
- Submission of Level 2 milestone "ASC Petascale Environment infrastructure Deployment Plan" planning document

Preliminary planned activities in FY09:

- File system preparation for Sequoia Petascale Environment
- Deploy prototype NFSv4 as NAS server between tri-lab sites
- Begin evaluation of parallel NFSv4.1 (pNFS)

WBS 1.5.4.5-LLNL-003 Networking and Test Beds

The LLNL Networking and Test Beds project provides research, performance testing, capability testing, and analysis for the file system, network, and interconnect subsystems in support of current and future systems and environments. This work relies heavily on an adequately provisioned test bed, skilled staff, and collaborations with vendors.

This project will test various hardware and software components to quantify the features, performance, reliability, security, and interoperability of the products and broader technology base. The information acquired as a result of this project will be used to help determine an integrated architecture and resultant procurements for these subsystems.

Planned activities in FY08:

- In support of future petaFLOPS systems, perform research and testing for technologies and products pertaining to interconnects, local area networks and WANs and NSA Type 1 encryptors, and file system servers, clients and disks, with special focus on large 10GigE switches, additional features in InfiniBand, other interconnect and 10GigE related technologies, and 10GigE NSA Type 1 encryptors
- Study mitigation of disaster affects, including off-site storage of critical data for recovery and back up network paths
- Apply testing results to optimize the functionality, performance, reliability, manageability, and security of the I/O services supporting these computing systems

Expected deliverables in FY08:

- Install and integrate switching infrastructure for TLCC07 and Sequoia ID (Dawn) system

Preliminary planned activities in FY09:

- Continue to leverage tri-lab activities in I/O related hardware and software, and seek to improve the reliability, performance and manageability of the I/O sub-systems in production
- Research and test to determine which technologies and products should be considered for insertion into production to meet the growing I/O performance and capacity requirements

WBS 1.5.4.5-LANL-001 File Systems and Input/Output Project

The File Systems and I/O Project provides end-to-end, high-performance networking and scalable I/O infrastructure for the ASC program. It also delivers high bandwidth, low-latency interconnect technologies for the ASC compute platforms. The ASC

program requires system and storage area network bandwidths at over 100 GB/sec., global file system I/O rates beyond 100 GB/sec., and latencies in the 1 microsecond range. All this performance must be provided in an integrated, usable, and secure way. Data transfer and storage bottlenecks are now a critical concern for current-generation, high-performance computing environments. Successfully meeting the ASC programmatic milestones requires carefully balanced environments in which the I/O infrastructure scales proportionally with increased ASC platform capabilities and application data needs.

This project is a coordination point for planning of all online storage, network, and data movement activities within the ASC program at LANL. These capabilities include online file systems such as the NFS complex and local supercomputer file systems, GPFS development, deployment and management, scalable I/O middleware development and support, interconnect technology development and deployment, and storage Area Networking development and deployment.

Planned activities in FY08:

- Continue to work on dead I/O node detection work for deployment
- Begin working on recovery from PaScalBB lane switch failure
- Begin working on recovery from storage node failure
- Work on deployment of I/O and storage infrastructure for Roadrunner Phase 3
- Begin testing NFS single name space service
- Assist users with I/O issues including help with the nearly 1-petaFLOP science runs
- Participate in writing of Level 2 petascale infrastructure milestone
- Detailed assessment and plan for petascale environment requirements

Expected deliverables in FY08:

- Deploy dead I/O node detect and reroute on all 3 GPFS environments red, yellow, turquoise
- Demonstrate initial PaScalBB switch and storage level fail over (full fail over end to end for I/O)
- Pending the October 2007 decision to proceed with Roadrunner phase 3, deploy I/O infrastructure to support this Petascale hybrid computing system
- Provide scalable NFS single name space service report to assist in decision for deployment
- Provide input to the Level 2 petascale infrastructure milestone

Preliminary planned activities in FY09:

- Deploy full I/O path fail over for GPFS
- Deploy global scalable NFS for all three environments (yellow, red, and turquoise)
- Assist users with I/O issues on Roadrunner phase 3, with special focus on weapons applications
- Deploy second generation scalable metadata operation

WBS 1.5.4.5-LANL-002 Archival Storage Design and Development

This project includes services for HPSS and PSI software development by LANL for the purpose of supporting ASC customers from LANL, LLNL, and SNL. These services include collecting user requirements for changes and upgrades to HPSS and PSI, developing plans for implementing user requirements into the codes performing the design and development work for upgrading the codes, and providing second-level support for the archive storage deployment team. The project works with the consulting office and archive storage deployment team to troubleshoot problems experienced with storing and retrieving data from the archive. The HPSS portion collaborates with tri-lab developers for implementing solutions that meet ASC requirements for all three labs. The PSI portion collaborates with LANL colleagues on user interface issues and ensures that PSI functions with each new release of HPSS.

Planned activities in FY08:

- Collect requirements for next releases of HPSS and PSI
- Create plans for implementing user requirements
- Design, develop, and test next versions of HPSS and PSI
- Work with customers in developing archive storage performance metrics
- Design, develop, and test HPSS release 8.x, including multiple meta-data servers and partitioned DB2 meta-data
- Provide second level support for archive storage deployment team
- Support for resource manager supported by PSI
- Provide HPSS user, installation, and administration manual updates
- Consider making access control lists portable and seamless across the tri-labs
- Evaluate file lifetimes for expiration of data within the archive
- Develop and implement formal and informal customer satisfaction metrics and measurement techniques

Expected deliverables in FY08:

- Deliver HPSS v7.1
- Provide new version of PSI with enhanced HTAR support
- Design, development, and implement HPSS/PSI resource manager
- Implement improved metrics reporting

Preliminary planned activities in FY09:

- Continued Level 2 support for the production archive
- HPSS 9.x requirements and design document
- Initiate next generation archive planning

WBS 1.5.4.5-SNL-001 Scalable Input/Output and Storage Systems

The Scalable I/O project provides scalable I/O infrastructure for ASC platforms in support of national security simulation needs and a mission-responsive computational environment. Continued R&D in parallel file systems and I/O libraries is critical to I/O for petascale and future exascale platforms.

HPSS provides flexible and scalable hierarchical storage management that supports a mission-responsive computational environment. HPSS uses cluster and storage area network technology to aggregate the capacity and performance of many computers, disks, and tape drives into a single virtual file system of exceptional size and versatility.

The Scalable I/O project works with vendors and researchers to provide reliable, high-performance, easily used scalable I/O libraries and file systems that make optimum use of disk I/O rates for petascale and beyond. It also educates users on how to ensure optimum I/O performance from their applications. The I/O library R&D contributes to the integration of current and future compute platforms, visualization rendering engines, and data management pre- and post-processor servers. The project works to ensure that higher-level I/O libraries effectively use the lower level libraries; in particular, HDF5, MPI-I/O, industry, and research file systems must work well together on all ASC platforms.

HPSS manages hundreds of terabytes to petabytes of data on disk and robotic tape libraries. It provides highly flexible and scalable hierarchical storage management that keeps recently used data on disk and less recently used data on tape. The single virtual file system approach enables HPSS to easily meet otherwise unachievable demands of total storage capacity, file sizes, data rates, and number of objects stored. HPSS provides a variety of user and file system interfaces ranging from ftp, samba, and nfs to higher performance pftp, client API, local file mover, and third party data transfer software.

Planned activities in FY08:

- Provide high-performance, reliable parallel file system I/O for Red Storm; includes SYSIO library that provides virtual file system capability for Red Storm compute nodes
- Optimize performance and test scalability of Lustre and intermediate I/O libraries (MPI-IO and HDF5) on Red Storm; work with application developers to optimize I/O performance on Red Storm and Purple
- Work with vendors and academia to further research in parallel file systems and I/O libraries for petascale and future exascale systems
- Enhance HPSS to improve small file performance and provide reliable, high-performance data transfer capabilities
- Lead I/O research for HPC systems (POSIX standards committee, parallel NFS, Northwestern University MPI-IO work, and Lightweight File System)

Expected deliverables in FY08:

- Provide I/O and storage expertise for the tri-lab petascale environment milestone
- Testing and optimization of the MPI-IO, parallel HDF5, and NetCDF libraries on Red Storm and the RoSE clusters; testing support for end-to-end I/O optimization
- Continue HPSS development to improve small file performance; provide reliable, high-performance data transfer capabilities; and provide Lightweight Directory Access Protocol (LDAP) support, becoming proficient in installation, configuration, and operation of HPSS IBM LDAP Server for systems in Kerberos environments

Preliminary planned activities in FY09:

- It is expected that most of the FY08 research and development of parallel file systems and I/O libraries will yield candidate technologies for future exascale systems; these technologies, coupled with advances in network transport services, such as RDMA and service guarantees, will create new opportunities to push state-of-the-art in I/O

WBS 1.5.4.5-SNL-002 Advanced Networking

Petascale computing platforms require connectivity to other petascale resources such as storage, visualization, and archive systems. A network that scales up to the performance of these platforms is required to maintain a balanced, efficient system. This project defines the requirements, technologies, and architectures that will scale the network up to the petascale range. Also, ASC large-scale Linux clusters increasingly use InfiniBand for their inter-process communication fabric. This project provides national leadership in InfiniBand's evolution as a HPC technology by influencing the development of the OpenFabrics/OpenIB common Linux software stack effort, as well as contributing to InfiniBand awareness in the OpenMPI messaging library.

The Advanced Networking project provides next-generation networking technologies in the 10–100-gigabit-per-second performance range spanning the range from system interconnect networks to the wide-area environment. Advanced capabilities are investigated to enhance ASC computational throughput and efficiency. Security, problem diagnostic capabilities, and performance monitoring of network services are also investigated and developed as part of this project. This work will provide the definition of local and remote network infrastructures to support a petascale computing environment and allow users to access remote ASC resources as if they were locally connected. New technologies and protocols will be tested and evaluated. The development laboratory and our analytical modeling capabilities will be utilized to support the design of the petascale computing environment.

Planned activities in FY08:

- Develop analytical techniques to increase the reliability and availability of the network environment; deficiencies will be identified and mitigation efforts will be proposed to ensure technologies will be available to support the petascale environment
- Manage ASC OpenFabrics development activities; optimize MPI collective performance for increased productivity; and extend SNL capabilities for predictive modeling of computer architectures
- Prototype and evaluate new implementations of Lustre file system, parallel NFS with RDMA, and emerging storage area network technology
- Support the tri-lab Level 2 petascale environment milestone

Expected deliverables in FY08:

- Develop a roadmap to deliver large scale switching infrastructures to support the petascale systems
- Develop and demonstrate tools that can provide automated monitoring and analysis of large-scale networking systems to provide the required redundancy and high availability
- Demonstrate large-scale solutions that provide highly efficient utilization of bandwidth from the processing cores through the network interface cards to take advantage of the large switching infrastructure
- Define FY09 OpenFabrics project plans for petascale systems

Preliminary planned activities in FY09:

- Implement OpenFabrics project plans for petascale systems

WBS 1.5.4.5-Y12-001 Applications in Support of Manufacturing Production and Connectivity

This project will support the utilization of ASC codes and computing resources to solve production manufacturing problems through modeling and simulation. The project will include support for connecting to ASC computing resources and job submission, execution and visualization. The project will also support the implementation of a compute cluster at Y-12 to provide the infrastructure necessary to test applications and scenarios before deployment on larger ASC resources. Development and implementation of software to support the solution of manufacturing problems will also be supported by the project. Finally, participation in NWC ASC-related activities will be covered.

The function of this project is to support the utilization of ASC codes and computing resources to solve production manufacturing problems through modeling and simulation. A cluster compute capability and software to meet Y-12 requirements will be implemented at Y-12.

Planned activities in FY08:

- Utilize Y-12 and remote ASC cluster resources for production manufacturing problems
- Develop plan to integrate Y-12 compute resource with ASC environment
- Participate in Nuclear Weapons Complex ASC activities

Expected deliverables in FY08:

- Utilization of Y-12 and remote ASC cluster resources for Production Manufacturing problems
- Develop plan to integrate Y-12 compute resource with ASC environment
- Participation in Nuclear Weapons Complex ASC activities
- Host Computational Manufacturing conference

Preliminary planned activities in FY09:

- Expansion of Y-12 cluster resource
- Development of software to support Y-12 requirements
- Participation in Nuclear Weapons Complex ASC activities
- Participate or host Computational Manufacturing conference

WBS 1.5.4.6: Pre- and Post-Processing Environments

This level 4 product provides integrated environments to support end-user simulation set up, and post-processing visualization, data analysis, and data management. The scope of this product includes planning, research, development, integration and deployment, continuing customer/product support, and quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include tools for optimized problem set-up and meshing, metadata and scientific data management, and general-purpose and application-specific visualization, analysis, and comparison. Research includes innovative data access methods and visualization of massive, complex data—the use of open-source foundations will continue to be an important strategy for development of shareable advanced techniques. The product

must develop solutions to address interactivity, scaling and tri-lab access for petascale platforms, and data analysis techniques needed to support effective verification and validation and comparative analysis. Solutions for emerging platform architectures may in turn require customization and/or re-architecting of software to leverage hardware features. A continuing emphasis will be placed on tools for improving end-user productivity. The product also provides and supports infrastructure including office and collaborative space visualization displays, mechanisms for image data delivery, and graphics rendering hardware.

Pre- and Post-Processing Environments Deliverables for FY08

- Deliver new comparative analysis techniques, feature extraction techniques, and synthetic diagnostics in VisIt
- Deliver new compression techniques in Silo and HDF5
- Resource-optimized transfers and operations for terabyte datasets in Hopper and Chopper
- Release general-purpose file system monitoring tools for GPFS, ZFS, and NetApps file systems
- Integration of the FY07 Powerwall drivers
- Production deployment of stereo in the Terascale Simulation Facility
- Video infrastructure upgrade for the LLNL Building 111 Visualization Work Center
- New weapons science visualizations for the TBI program
- Classified papers on weapons science topics jointly co-authored with X Division designers to be presented at the NEDPC 2008 conference
- Classified published journal articles jointly co-authored with X Division designers for publication in *Defense Research Review*
- Other visualizations and data analysis products to address a variety of DSW-related activities including the closing of outstanding SFIs
- Scientific visualization support for briefings
- Develop workflow infrastructure to automate comparative and quantitative visual data analysis using the production visualization tools; study improving the comparative visual analysis process by gathering and using knowledge gained from using a workflow infrastructure
- Investigate and document approaches to interfacing new data analysis methods into the production visualization tools
- Design improved data compression schemes for distance visualization. Incorporate them into level-of-detail streaming visual analysis infrastructure
- Investigate and document approaches to accelerating visualization, analysis and rendering kernels on the Cell processor
- Complete petascale Level 2 milestone based on information gained by prototyping new comparative visualization and scaling approaches
- Tri-lab FY08 milestone: petascale environment document
- Advanced V&V capabilities in support of high energy density physics (HEDP), including vector field support and delivery of vector visualization techniques

- Semi-supervised learning and active learning algorithms in AVATAR Tools; develop and implement in FCDMF (feature characterization) techniques for "cross-model" comparison of characterizations
- Completed deployment of unclassified data analysis infrastructure; initial upgrades toward alignment with upgraded Red Storm capability

WBS 1.5.4.6-LLNL-001 Scientific Visualization

This project provides high-end visualization and analysis capabilities for all ASC simulation codes at LLNL, broken down into three activities: 1) research into new analysis and visualization techniques, 2) research into new compression techniques, and 3) deployment of the new techniques into the analysis and visualization tool VisIt, scientific I/O libraries used by ASC, custom applications, or new libraries.

The team members of this project have extensive experience handling data processing needs of the Weapons & Complex Integration Principal Directorate (formerly called DNT, Defense & Nuclear Technologies) users. Our researchers are world-renowned for their expertise with data and image compression, including on-the-fly compression of very large datasets; work on multi-resolution streaming architectures for data processing and visualization; topological “summary” techniques for scientific datasets; and view-dependent rendering and data management techniques. We have extensive experience with parallel rendering techniques, both for tiled displays and compositing architectures. The team exploits the latest capabilities of clustering hardware, graphics processing unit advances, and parallel storage systems.

Planned activities in FY08:

- Maintain and enhance the existing suite of tools and libraries developed by the visualization project
- Perform research and development in topology and data compression
- Begin to investigate methods to leverage multi-core technology
- Participate in planning activities for petascale data analysis capabilities

Expected deliverables in FY08:

- Deliver new comparative analysis techniques, feature extraction techniques, and synthetic diagnostics in VisIt
- Deliver new compression techniques in Silo and HDF5

Preliminary planned activities in FY09:

- Perform research and development in topology and level of detail
- Continue to investigate methods to leverage multi-core technology

WBS 1.5.4.6-LLNL-002 Scientific Data Management

This project is focused on providing users with better ways to generate, organize, search, access, extract, compare, track, and archive large-scale scientific data. This is achieved through two major themes: production applications for enhancing data access and organization, and data discovery techniques and tools for representing, exploring and extracting pertinent information. Data discovery develops scalable algorithms for the interactive exploration of large, complex, multi-dimensional scientific data. It applies and extends ideas from data mining and pattern recognition in order to improve the way in which scientists extract useful information from data.

The Scientific Data Management tools team has expertise in data transport protocols, graphical user interfaces, web technologies, data representation, databases, and advanced system architectures. The team has decades of combined experience in producing heavily-used production applications. The data discovery effort focuses on the Sapphire data mining project, which has expertise in data mining, image analysis, video tracking, statistical techniques, and pattern recognition. The Sapphire team's efforts have garnered many software patents and an R&D 100 award.

Planned activities in FY08:

- Maintain and enhance the existing suite of tools and libraries developed by the scientific data management project
- Add support for tri-lab data movement protocols to facilitate Hopper being part of the Tripod software stack
- Extend the Lustre Monitoring Tools collection and display applications
- Analyze high-fidelity simulations of Rayleigh-Taylor instability and characterize experimental images for DNT physicists and provide code surrogates and identification of key features for DNT simulation data

Expected deliverables in FY08:

- Resource-optimized terabyte transfers and operations for terabyte datasets in Hopper and Chopper
- Release general-purpose file system monitoring tools for GPFS, ZFS, and NetApps file systems

Preliminary planned activities in FY09:

- Continue to maintain and enhance the SDM suite of tools and libraries
- Investigate tighter coupling of Hopper user interface with local desktop
- Analysis of simulation and experimental data as requested by DNT physicists

WBS 1.5.4.6-LLNL-003 Systems and Services Integration

This project is focused on visualization servers, displays, and operational support of the pre- and post-processing product resources for ASC computing customers. The server effort includes planning, test bed prototypes, component and system testing, computer security and system deployment, system administration, and maintenance of visualization servers. The display effort includes high resolution / performance display devices for user offices and collaborative use areas. Operational support includes managing the use of Powerwall facilities and the associated servers, running video production labs, and applying and consulting on software including resource management tools, movie players, animation and visualization packages.

The project has expertise in systems administration, software development, a wide range of visualization hardware and software tools and techniques, advanced displays and related infrastructure, windowing systems, graphics processing units, video production, computer animation and 3D modeling, and user interfaces. Hardware capabilities include five production visualization servers with a variety of central processing units, interconnects, graphics cards, and disks. The Linux clusters are connected to a Lustre file system and video display infrastructure that drives Powerwalls and smaller displays. We install, maintain, and consult on a wide variety of software tools. The project supports demonstrations on the Powerwall, and we have the capability to visualize simulation data and create movies to be shown on Powerwalls or

laptops. The project maintains both an unclassified and classified video production lab, each of which includes desktop systems with video editing software, 3D modeling and animation tools, and an assortment of video peripherals. We create DVDs and videotapes as needed in addition to supporting live demos.

Planned activities in FY08:

- Maintain all of the visualization servers and advanced displays currently supported, and to provide operational support for all visualization facilities
- Provide consulting and support for graphics and video production, and we will keep the production graphics software environment current and consistent
- Future hardware evaluation, technology watch and test beds

Expected deliverables in FY08:

- Integration of the FY07 Powerwall drivers
- Production deployment of stereo in the Terascale Simulation Facility
- Video infrastructure upgrade for the Building 111 Visualization Work Center

Preliminary planned activities in FY09:

- Maintain all of the visualization servers and advanced displays currently supported, and to provide operational support for all visualization facilities
- Provide consulting and support for graphics and video production, and we will keep the production graphics software environment current and consistent
- Future hardware evaluation, technology watch, and test beds

1.5.4.6-LANL-001 Visualization and Insight for Petascale Simulations Project

This project is comprised of two parts: "improved/automated data analysis capabilities" and "scaling to petascale interactive visualization and analysis."

The improved/automated data analysis capabilities project encompasses ongoing work on comparative and quantitative visual analysis of ensembles of data and the automation of this process. Specifically, we are using workflow systems to support this approach. When dealing with the complexity of ensembles of data, multi-variables, and massive datasets it is important to automate as much of the analysis process as possible. One benefit of following this approach is since workflow systems are script-based, the work is directly applicable to integrating production tools with each other and with prototype visualization tools to provide access to best functionality of each of these tools.

The scaling to petascale interactive visualization and analysis project encompasses identifying appropriate hardware resources to support petascale visualization, working on data reduction based software techniques such as intelligent data streaming to reduce the volume of data need to be moved off the petascale platform, and exploring the use of visualization kernels that run on the petascale platform (including data analysis, visualization, and rendering methods). Running on the petascale platform is important for two reasons: 1) to get around the bandwidth limitations imposed by the architecture (for example if you can analyze as the data is computed, you can avoiding writing to disk and then re-reading the data); and 2) for performance, using the power of the petascale platform to analysis petascale data.

The Petascale Visualization and Data Analysis milestone will be managed out of this project.

Planned activities in FY08:

- Improved and automated data analysis capabilities to support improved accuracy and fidelity in physics
- Scaling to petascale interactive visualization and analysis

Expected deliverables in FY08:

- Develop workflow infrastructure to automate comparative and quantitative visual data analysis using the production visualization tools; study improving the comparative visual analysis process by gathering and using knowledge gained from using a workflow infrastructure
- Investigate and document approaches to interfacing new data analysis methods into the production visualization tools
- Design improved data compression schemes for distance visualization; incorporate them into level-of-detail streaming visual analysis infrastructure
- Investigate and document approaches to accelerating visualization, analysis and rendering kernels on the Cell processor
- Complete petascale Level 2 milestone based on information gained by prototyping new comparative visualization and scaling approaches

1.5.4.6-LANL-002 Production Systems for Visualization and Insight Project

The primary goal of this project is to provide world-class visualization infrastructure and support services in an efficient and cost-effective manner through six sigma processes.

Visualization and visual analysis are essential tools needed by code teams and designers in understanding the terabytes of data that are generated in a single simulation run. This project, also referred to as "Production Visualization," provides visualization services from the machine to the desktop for users in the ASC program. Most importantly, people funded by Production Visualization work with code teams and designers to visualize their datasets, train them so that they can visualize their own datasets, assist them in using the large facilities, and assist in giving briefings in the facilities to high-level visitors. Software currently supported at full strength by the Production Visualization Project includes EnSight, EnLiten, EnVideo, Vista, PoP, and GMV. This project also supports and maintains LANL's large visualization facilities, including the CAVE, the Powerwall Theater, and the co-laboratories. It maintains the visualization infrastructure, which delivers video from the machines to the users' desktops. Finally, it assists in the process of bringing new machines up by troubleshooting graphics systems on these machines, by performing the visualization software integration tasks needed, and by installing and maintaining critical visualization software on ASC machines.

Capabilities include:

- Custom visualizations for Code Teams A and B, and ESA Division
- Development and support of PoP and GMV in a conventional and petascale, hybrid environment

- Support of the large facilities, such the CAVE and the PWT, including assistance, training, and developing new tools to work with these facilities
- Briefing support
- Deployment and maintenance of any needed Visualization Corridor software

The goals of the Production Visualization project are six fold:

- To support and maintain the large visualization facilities and conduct briefings.
- To better integrate our visualization infrastructure into our mainstream computing production environment. This includes of our high-speed network and scalable GPFS.
- To support and maintain the PoP and GMV packages. To attain these goals, we plan to provide visualization, software, and facilities support at the same high quality as in the past and push the further integration of visualization resources into the general computing infrastructure. If the Roadrunner option is exercised, we would like to exploit the visualization -specific Roadrunner cell technology to enhance our visualizations both on the desktop and in the facilities.
- To support users in providing visualizations of simulations, experimental data and engineering CAD data on desktops and in facilities.
- To apply expert knowledge of weapons physics and the visualization software environment to iteratively develop solutions of interest to the program.
- To work directly with designers in physics-based, iterative discovery process using Petascale Visualization and Data Analysis enabled tool (Ensight)

Planned activities in FY08:

- Continue the support and maintenance of the large visualization facilities
- Conduct briefings as needed
- Support and maintain the PoP and GMV packages
- Support users in providing visualizations of simulations, experimental data and engineering CAD data on desktops and in facilities

Expected deliverables in FY08:

- High availability in the large visualization facilities
- Scientific visualization support for briefings
- PoP and GMV upgrades
- User visualizations

Preliminary planned activities in FY09:

- Continued support and maintenance of the large visualization facilities
- Continued support of users in providing visualizations of simulations, experimental data and engineering CAD data on desktops and in facilities

1.5.4.6-LANL-003 Physics-Based Simulation Analysis Project

The purpose of this project is to help LANL weapons designers to utilize the full power of the hardware and software infrastructure for visualization and data analysis developed and deployed by ASC to improve the physics understanding of their

weapons simulations. To achieve this goal the project has deployed within the design community in X Division a small group of individuals with expert knowledge in both visualization and weapons science to work directly with the designers. The job of this small group of experts is to help designers apply the full potential of the ASC visualization and analysis infrastructure to solve their analysis problems and to promote new weapons science discoveries using the ASC codes.

This small group of individuals has played a major role in the recent successes of the TBI program at LANL. Their work in combining ASC visualization with the new capabilities provided by the ASC codes has been one of the major factors responsible for the new discoveries being made by the TBI program.

In addition to working directly with the design community on its visualization and analysis problems, this group is also responsible for all LANL activities related to the EnSight visualization and data analysis software. This includes maintaining the EnSight software installation Lab-wide, providing local user support in the use of the software and acting as a bridge between the LANL design community and the EnSight developers at Computational Engineering International for problem reporting and resolution and for new feature requests. This group also directs all subcontracts that LANL has with Computational Engineering International related to new EnSight development and to on-site training and consulting.

At the moment the project has three such experts working directly with X Division designers on applications of visualization and analysis to the solution of weapons program problems. Each of these individuals has a physics or nuclear engineering background with dual expertise in both ASC visualization and weapons science. All have extensive experience with the ASC visualization infrastructure developed by LANL.

Planned activities in FY08:

- Application of expert knowledge of weapons physics and the visualization software environment to iteratively develop solutions of interest to the program
- Work directly with designers in physics-based, iterative discovery process using Petascale Visualization and Data Analysis enabled tool (EnSight)
- Support and maintain the EnSight software and direct EnSight development activities under the new LANL EnSight development contract to Computational Engineering International
- Document our work with joint publications co-authored with X Division designers on a variety of weapons science topics

Expected deliverables in FY08:

- New weapons science visualizations for the TBI program
- Other visualizations and data analysis products to address a variety of DSW-related activities including the closing of outstanding SFIs

Preliminary planned activities in FY09:

- Continue to promote new discoveries in weapons science by advanced applications of visualization and data analysis in programs such as TBI
- Continue to document the results of our activities with classified papers and publications on weapons science topics jointly co-authored with X Division designers

WBS 1.5.4.6-SNL-001 Scalable Data Analysis

The Scalable Data Analysis project enables ASC customers to analyze, explore, and understand complex data results within CSSE's responsive design and analysis capability. A prominent theme of near term work is support of verification and validation (V&V) requirements. R&D in this project addresses a spectrum of needs, from tools promoting comparison and analysis of large numbers of runs to support for advanced analytic components needed by ASC simulations.

Ideally, analysis should not be impacted by data size; query results, images, and other analysis products should be presented quickly to the analyst, to promote deep thinking about results. New algorithms, presentation techniques, and data characterizations are needed to achieve this ideal. In addition, as data size increases, we must integrate information analysis with scientific visualization to truly understand our results.

This project is focused on providing advanced customer-centered capabilities within an open source production framework (ParaView). Our scalable tools allow investigation of data on a variety of platforms—everything from a laptop to a cluster. This allows users to interact with their data, whether it fits on a PC or is located on a remote cluster in another state. In addition, our tools are beginning to deliver advanced analysis capabilities, in addition to traditional scientific visualization, that promote understanding of large data and investigation of ensembles of runs necessary for ASC's V&V goals.

Planned activities in FY08:

- Support of FY08 tri-lab milestone: Infrastructure deployment plan for ASC petascale environments
- Continued collaboration with SNL V&V program element and targeted customers to deliver V&V capabilities with broad programmatic impact; supports FY08 milestone: "Deliver post-processing tools that enable verification and validation of FY08 HEDP simulations"
- R&D of ensembles of views for V&V and advanced analysis of large data, ensembles of runs, and comparison of runs
- Continue advances in comparison and characterization techniques for large data

Expected deliverables in FY08:

- Tri-lab FY08 milestone: Infrastructure Deployment Plan Document
- Advanced V&V capabilities in support of HEDP, including vector field support and delivery of vector visualization techniques
- Simplified python scripting interface client to distributed analysis engine
- Ongoing releases of SNL-specific R&D in ParaView 3.x
- Semi-supervised learning and active learning algorithms in AVATAR Tools; develop and implement in FCDMF techniques for "cross-model" comparison of characterizations

Planned activities in FY09:

- Initial release of advanced algorithms/architectures for data mining scientific data
- Advanced characterization tools for comparing large scale simulated and experimental data

- Advanced ensembles of views, scripting and algorithms for V&V and analysis of large data

WBS 1.5.4.6-SNL-002 Visualization Deployment and Support

This project deploys and supports software and infrastructure to deliver production data analysis and visualization capabilities. This project works closely with the Scalable Data Analysis project to provide a bridge from development to production; with many ASC projects to deliver its tools as part of an integrated modeling and simulation capability, including tri-lab; and with ASC platform and other infrastructure deployment activities to deliver a complete end-to-end system environment. This project also interfaces with applications and users to refine and enable use of ASC capabilities.

The project provides installation, testing, benchmarking, and end-user support for post-processing tools and utilities. Installation and support span the diversity of platforms in the distributed tri-lab environment. The project works with the other ASC labs to ensure that expected tools are deployed as needed to support the tri-lab computing environment. As to infrastructure, the project provides scalable data analysis systems that interactively analyze, visualize, store, and archive output from ASC computers.

Planned activities in FY08:

- In partnership with Scalable Data Analysis project, support delivery of post-processing tools that enable verification and validation of HEDP simulations
- Provide post-processing support for UQ application demonstration of design-through-analysis environment
- Support infrastructure planning and evolutionary deployment toward petascale environments
- Ongoing incremental deployment, upgrade, and refresh of data analysis system infrastructure; integration of infrastructure with new tri-lab TLCC systems

Expected deliverables in FY08:

- New release(s) of ParaView containing emerging V&V data analysis features
- Completed deployment of unclassified data analysis infrastructure; initial upgrades toward alignment with upgraded Red Storm capability

Preliminary planned activities in FY09:

- Incremental deployment of capabilities toward petascale environments
- Ongoing delivery and support of up-to-date post-processing capabilities within an integrated responsive design and analysis capability.

WBS 1.5.5: Facility Operations and User Support

This sub-program provides both necessary physical facility and operational support for reliable production computing and storage environments as well as a suite of user services for effective use of ASC tri-lab computing resources. The scope of the facility operations includes planning, integration and deployment, continuing product support, software license and maintenance fees, procurement of operational equipment and media, quality and reliability activities, and collaborations. FOUS also covers physical space, power and other utility infrastructure, and LAN/WAN networking for local and remote access, as well as requisite system administration, cyber-security, and operations services for ongoing support and addressing system problems. Industrial and academic collaborations are an important part of this sub-program.

WBS 1.5.5.1: Facilities, Operations, and Communications

This level 4 product provides necessary physical facility and operational support for reliable production computing and storage environments. The scope of this product includes planning, integration and deployment, continuing product support, software license and maintenance fees, procurement of operational equipment and media, quality and reliability activities and collaborations. This product also covers physical space, power and other utility infrastructure, and local area network/wide area network (LAN/WAN) networking for local and remote access, as well as requisite system administration, cyber-security and operations services for ongoing support and addressing system problems.

Facilities, Operations, and Communications Deliverables for FY08

- Deploy TLCC07 systems at LLNL, LANL, and SNL
- Prepare site and infrastructure for the Sequoia ID system
- Develop and deploy a new identity/account management system
- Deploy the next generation File Interchange System
- Enhance core networks to support hundreds of full bandwidth 10 Gigabit Ethernet (GigE) connections for TLCC and 1000s for Sequoia
- Deploy and operate full production of Roadrunner Base System known as Redtail in secure network and Yellowrail on the open protected network
- Prepare LANL SCC machine room for Roadrunner Phase 3
- Improve existing 10 Gigabit network infrastructure to support production systems
- Deliver classified and unclassified capability and capacity computing cycles for the tri-lab, with unclassified time for the Academic Strategic Alliance Program partners
- Deliver a tri-lab common software environment on the TLCC system
- Deliver a reliable monitoring system for the networking infrastructure connecting ASC tri-lab classified computing environments

WBS 1.5.5.1-LLNL-001 System Administration and Operations

The System Administration and Operations project provides for the ongoing system administration and computer operations functions for the successful management and support of the ASC platforms and computing environment.

Capabilities include highly skilled system administration to ensure installation, integration, and ongoing support of ASC platforms including operating system and software configuration; feature, functionality, and security patches; troubleshooting, analysis, and diagnosis. This project also includes a 24/7 operational monitoring capability for unclassified and classified computing environments consisting of large-scale computing platforms, infrastructure components, and networks.

Planned activities in FY08:

- Preparations for the Sequoia ID system
- Preparations for and deployment of LLNL SUs from the TLCC procurement
- Ongoing support for Purple, BlueGene/L, and other capacity systems
- Retirement of Lilac and ALC

Expected deliverables in FY08:

- Deployment of TLCC SUs
- Ongoing support for Purple, BlueGene/L, and other capacity systems

Preliminary planned activities in FY09:

- Ongoing support for Purple, BlueGene/L, and other capacity systems
- Continued preparations for and deployment of the Sequoia ID system
- Preparations for and deployment of LLNL SUs from the TLCC procurement

WBS 1.5.5.1-LLNL-002 Software and Hardware Maintenance, Licenses, and Contracts

This project provides for vendor-provided hardware and software maintenance, support, licenses, and development contracts. For laboratory-maintained systems, the project provides hardware maintenance capabilities including component inventory and replacement.

Capabilities include negotiated hardware and software maintenance and license contracts to ensure a robust ASC computing environment and to protect the computational investment of the NNSA. Targeted development contracts to enhance the capabilities of specific software components are also included.

Planned activities in FY08:

- Track and place contracts and licenses needed for system operations and vendor support
- Self-maintenance of designated systems

Expected deliverables in FY08:

- Contracts and licenses needed for system operations and vendor support
- Self-maintenance of designated systems

WBS 1.5.5.1-LLNL-003 Computing Environment Security and Infrastructure

The Computing Environment Security and Infrastructure project provides for computing environment infrastructure services, software, servers, workstations, and desktop systems necessary for the efficient, effective, and secure operation and support of large-scale ASC platforms.

Capabilities include ongoing integration, development, and support on unclassified and classified networks of robust infrastructure environments to support large-scale ASC platforms including but not limited to name and time services, backups, staff productivity tools (for example, e-mail and messaging), cyber security tools and technologies.

Planned activities in FY08:

- Develop and deploy an identity/account management system
- Deploy the next generation file interchange system
- Ongoing security environment support activities

Expected deliverables in FY08:

- Deploy an identity/account management system
- Deploy the next generation file interchange system

Preliminary planned activities in FY09:

- Continued investigation and analysis of new technologies for upgrade or replacement of existing technologies
- Ongoing security environment support activities

WBS 1.5.5.1-LLNL-004 Facilities Infrastructure and Power

The Facilities Infrastructure and Power project provides for the necessary physical facilities, utilities, and power capabilities to support staff and the ASC computing environment.

Capabilities include adequate raised floor space, cooling facilities, and power to site large-scale ASC platforms. In addition, needed office, meeting room, and auxiliary space to enable a highly motivated and effective staff is part of this project.

Planned activities in FY08:

- Initiate and track the progress of the institution components of the 15-megawatt power upgrade for Building 453 for 2011 completion
- Ongoing maintenance and support of existing computational and staff facilities
- Continuing analysis of future modifications and/or expansion of facilities that will be needed by future ASC systems

Expected deliverables in FY08:

- Site preparations for the Sequoia ID system and TLCC systems

Preliminary planned activities in FY09:

- Continue the 15-megawatt power upgrade progress tracking
- Ongoing maintenance and support of existing computational and staff facilities

- Continuing analysis of future modifications and/or expansion of facilities that will be needed by future ASC systems

WBS 1.5.5.1-LLNL-005 Classified and Unclassified Facility Networks

The Classified and Unclassified Facility Networks project provides the architecture design, planning, procurement, deployment, and operational support of the classified and unclassified facility networks.

Capabilities include a thorough understanding of the resource deployment roadmap is acquired by participating in ongoing facility-wide planning efforts that include the archival storage, visualization, platforms, capacity computing, and file systems. Network design, procurements, and deployments are updated and scheduled to accommodate these plans and ensure the network connectivity, performance, reliability, security, and operational support is available for the facilities to meet the requirements of all subsystems is also part of this project.

Planned activities in FY08:

- Network connections to support the Sequoia ID and TLCC systems
- Enhance core networks to support hundreds of full-bandwidth 10-gigabit-Ethernet (GigE) connections for TLCC and thousands for Sequoia
- Procure needed hardware and upgrade the Open Lab Net to 10 GigE
- Evaluate the use of other alternative technologies for use as interconnect for future systems
- Ongoing local network performance, reliability, security, and operational support

Expected deliverables in FY08:

- Network connections to support the Sequoia ID and TLCC systems
- Deployment of the -gigabit Ethernet Open Lab Net connection

Preliminary planned activities in FY09:

- Network testing for Sequoia ID system and any new connections for TLCC systems
- Ongoing local network performance, reliability, security and operational support

WBS 1.5.5.1-LLNL-006 Wide-Area Classified Networks

The Wide-Area Classified Networks project provides the architecture design, planning, procurement, deployment and operational support of the classified wide-area networks, namely the DisCom WAN and the SecureNet WAN.

Capabilities include ongoing discussions with the tri-lab community are critical to this project to ensure that the network requirements for those remote users and facilities are mutually agreed upon and understood. This project must also plan far in advance to ensure the required NSA Type 1 encryption products are available, since these products are not commercial and have a long R&D and product development lead time. Operational support of these WANs also requires effective and regular communication with and cooperation between the tri-lab network support teams. These activities will help ensure the proper planning occurs for the WANs, and the operational support is effective for the broader tri-lab user community.

Planned activities in FY08:

- Track development and availability of 10-gigabit capable encryption units; procure and deploy for the classified network
- Continued operational support of the WANs and tri-lab connections

Expected deliverables in FY08:

- Procure and deploy a 10-gigabit encryption unit for classified network

Preliminary planned activities in FY09:

- Track requirements for WAN bandwidth as ASC systems approaching a petaFLOPS are deployed for tri-lab use
- Continued operational support of the WANs and tri-lab connections

WBS 1.5.5.1-LANL-001 High Performance Computing Operations Monitoring and Requirements Planning

This project covers the planning activities for computing operations, collection, and statistical evaluation of user requirements for computing resources and development of new metrics and data collection.

The primary capability of this function is to collect and understand user requirements for production computing resources and quality of service, and to develop new metrics and data collection and analysis techniques to assist these purposes.

Planned activities in FY08:

- Extending the infrastructure to collect useful and relevant application level usage data, including computing sub-system (CPU, memory, file system) usage data, and on reporting collected data in a timely and planning-productive manner.
- The infrastructure for collection of application-level data will merge into current infrastructure efforts for the collection and analysis of system-level RAS and operations correctness data, forming the HPC Monitoring Infrastructure Project, an activity integrating application-level and system-level RAS data to form a more complete view of RAS from an end-user perspective. This effort will be coordinated with the other nuclear weapons complex laboratories.

Expected deliverables in FY08:

- Completion of system-level RAS monitoring infrastructure
- Design of the data organization and interface to view integrated application-level and system RAS-level correctness and completion data
- Implementation of data transfer/organization of the application-level data to the RAS monitoring infrastructure
- Implementation of the data organization and interface for viewing the integrated data
- Deployment on one or more LANL HPC clusters

Preliminary planned activities in FY09:

- Continue to contribute to improvements in the quality and reliability of computing requirements collection and analysis, both at LANL and in the tri-lab arena

WBS 1.5.5.1-LANL-002 Roadrunner Phase 3 Delivery and Initial Deployment

This scope of this project is to take delivery and start the deployment of the Roadrunner Phase 3 System. It assumes a decision was made in October 2007 to proceed with the full procurement of this system.

Planned activities in FY08:

- Work with the vendor to ensure successful system testing at their assembly area prior to shipping the CUs to LANL. The CUs to be delivered to the Nicolas C. Metropolis Center for Modeling and Simulation (SCC) will be initially configured in the yellow unclassified network in preparation for the acceptance testing by the vendor. This system is architected to form the basis for an advanced architecture computing platform providing computing performance in the petaFLOPS realm.
- Prepare for delivery of Roadrunner Phase 3 CUs, which are expected at LANL before the end of FY08.
- Start evaluation of delivered system with diagnostic testing. Perform system tests on the overall compute system, and start the integration of the system into the LANL computing environment for production work.

Expected deliverables in FY08:

- Complete the system delivery of the Roadrunner Phase 3 System
- Start initial deployment of the delivered Roadrunner Phase 3 CUs

Preliminary planned activities in FY09:

- Directly involve management and members of the production computing staff on the integration aspects of the system. The system integration team will form a peer relationship with the vendor staff to implement and schedule the technical tasks identified in the Statement of Work. The production computing team will form a customer/vendor relationship with the vendor to provide an integrated support structure for the system for all operational issues. The areas currently identified for coordination include: resource management, near-line storage, HPSS/PSI, user training, user documentation, security plans, user environments/tools, account management, consulting readiness, and off-hours support.

WBS 1.5.5.1-LANL-003 Ongoing Network Operations

Ongoing Network Operations encompass all aspects of the network infrastructure that support ASC computing in the classified and, to some extent, the unclassified arena. This includes directly attached networks to HPC systems (machine area network), network backbones, the users' local area network, and the high-end WAN (DisCom WAN) connecting the tri-labs.

The network aspect of production computing provides the technical expertise and operation required to support high-end production computing and data storage services for the use of ASC computing. FY08 commitments include network support for ASC Roadrunner deployment and production and continued support for Lightning, Panasas global file system, and HPSS and access, via the DisCom WAN. These efforts also include network services in the classified partition, which are required to promote security while maintaining an environment that is acceptable to all. Increased network backbone bandwidth within the LANL campus is being implemented to support growth in HPSS storage from local and remote ASC computing and improve data delivery for

scientific visualization. Increased bandwidth over the DisCom WAN is being implemented for anticipated requirements from LANL users of the new remote systems and SNL and LLNL.

Planned activities in FY08:

- Network support: Designing, developing, acquiring, deploying, and supporting classified network hardware and services to support computational systems. This includes operating and maintaining services such as e-mail, authentication, and Web servers, plus operating and managing the high-performance network backbone, services networks, local paths to ESNet, SecureNet, DisCom and the WAN, the high-performance parallel interface/gigabit and ten gigabit networks.
- WAN operation and support: In close cooperation with LLNL and SNL, operate and maintain the recently upgraded DisCom WAN. Deploy new technologies to increase utilization, bandwidth, and reliability where warranted. Deploy new internet protocol encryptors to position WAN growth and control costs. Plan for bandwidth growth from two to four times current bandwidth if use by LANL users warrants such growth.
- Work closely with the Integrated Cyber Security Initiative project to deploy the new Enterprise Secure Network (ESN), which will not only include SecureNet operations but will also address directory, email, and other network services currently running locally at each site or over SecureNet. Plan for a demonstration of the new ESN working with an ASC application in the future when ESN is deployed.

Expected deliverables in FY08:

- Ongoing network support for all ASC systems at LANL is the deliverable for this task. Specific network upgrade efforts are deliverables in the next WBS sub-item.

Preliminary planned activities in FY09:

- Continued operation and maintenance of the LAN/MAN/WAN infrastructure. We also anticipate more integration with the ESN deployment. We expect further deployment and tuning of the I/O infrastructure to support Roadrunner Phase 3 as it increases in capability.

WBS 1.5.5.1-LANL-004 Network Infrastructure Integration

Infrastructure Integration encompasses the entire network infrastructure upgrades needed to support integration of new HPC computing platforms planned for FY08 and beyond. This includes designing, procuring, prototyping, testing, and installing network switches and routers needed to support the new HPC platforms and provide increasing bandwidth to HPSS, visualization systems, and customer workstations. In addition, this will also include integration of the network components into the HPC computing platforms and storage systems. This project will also include integrating new switches/routers into our network monitoring platforms and training operational personnel to monitor and troubleshoot the new systems.

Planned activities in FY08:

- Design and start integration of the network infrastructure to support the new Roadrunner Phase 3 system and its integration into the Integrated Computing Network infrastructure.
- Design and integrate the network infrastructure to support the new TLCC clusters.

- Extend our GPFS networks, PaScalBB, to support new HPC clusters and I/O links in both the Yellow and Red networks. We will also upgrade network infrastructure, where necessary, to increase bandwidth to HPSS, visualization platforms, and customer workstations.
- Work with system developers to stand up a new DISCOM landing platform to replace SGI systems in order to continue support of high-speed data transfers between LLNL/SNL and LANL:
- Provide support for system interconnect reliability and performance management.

Expected deliverables in FY08:

- Deploy new DisCom gateway, which will optimize use of the available DisCom bandwidth for parallel file transfers between Laboratories.
- Deploy network enhancements for new TLCCs.
- Complete design of I/O Network Integration for Roadrunner Phase 3.

Preliminary planned activities in FY09:

- Complete the integration of the Roadrunner Phase 3 system into the network infrastructure in preparation for full-scale operation.
- Involvement in developing and evaluating a petascale computing environment with emphasis on the network technology that will enable this environment envisioned at LANL in FY09 and beyond.

WBS 1.5.5.1-LANL-005 Ongoing Systems Operations

This area includes all services for systems operated by LANL for the purpose of developing and running user codes. These services include system configuration, computer security, resource management, system administration, and system operation and monitoring. The project works with users to troubleshoot problems experienced with running their applications on new systems, and also works with users to plan and carry out transitions off of older platforms.

Planned activities in FY08:

- System support (on-going): Conduct daily system administration with continuous monitoring of production systems and infrastructure servers
- Scheduling environment daily management by providing daily administration, monitoring, and problem resolution of software subsystems
- RAS (on-going): Continuously improve the end-to-end level of service as seen by the users. Conduct ongoing studies and improvement projects in the stability of large, integrated systems, including the development of improved diagnostic and monitoring capabilities
- RAS metric collection, reporting, analysis, and management (FY08 effort): FY08 RAS efforts build on FY07 design and initial deployment efforts for HPC systems
- System Operations: Ongoing operations require the staff to provide around-the-clock operations and monitoring of the scientific computing resources. This includes an increased level of system self-maintenance by this staff for various computing and storage systems
- Transition of Roadrunner base capacity system into full production
- Provide support for installation and integration of Roadrunner Phase 3

- Install TLCC for production and support them in production environment
- Continue to study the feasibility of reduced staff ("lights out") computer operations
- Implement additional formal and informal customer satisfaction metrics and measurement techniques
- Expand self maintenance to additional systems and components as practical
- Data storage operations for GPFS, NFS server complex, and archival storage (HPSS)

Expected deliverables in FY08:

- Decommission of QB, CA, CB, CC, and QSC computers
- Installation and deployment of TLCC systems with Tripod software stack
- Installation of Roadrunner Phase 3 system

Preliminary planned activities in FY09:

- Provide support for Roadrunner Phase 3 for at-scale Weapons Science calculations and for advanced architecture studies with integrated weapons applications.
- Pending results of siting decision, prepare for installation of next production tri-lab capability platform.
- Prepare additional (TLCC) capacity systems for production and support them in production environment
- Continue to provide system support, operations, and management for current systems
- Implement a reduced staff ("lights out") computer operations model if practical
- Implement improved RAS reporting
- Deploy a more extensive use of Tripod tools and capability on the TLCC system

WBS 1.5.5.1-LANL-006 Systems and Network Maintenance

This project obtains hardware and software support services for all production ASC systems and network assets operated by LANL. Support services include vendor hardware and software support contracts; contracts for parts sparing and logistics for self maintained equipment; software licenses; annual software maintenance and support fees; and contracts for telephone support and other remote hardware and software support requirements.

This project is a coordination point for all maintenance contracts for system and network hardware and software. Customers include ongoing operations and integration teams, projects and groups. Interfaces include third party vendors and suppliers and the LANL supply chain and accounting organizations.

Planned activities in FY08:

- Refining order entry and tracking capabilities including interfaces with customers, supply chain, program management, and accounting
- Forging even stronger relationships and understandings with customers and external interfaces
- Developing methods and procedures to ensure the capture of maintenance requirements for all new production resources well in advance of their required date

- Procuring and ensuring fulfillment of all maintenance services required by existing and new systems in FY08
- Planning for all maintenance services required in FY09 and beyond

Expected deliverables in FY08:

- New maintenance contracts for all infrastructure equipment procured for support of Roadrunner Phase 3 and TLCC system

Preliminary planned activities in FY09:

- Continued enhancements to contract for systems supporting Roadrunner Phase 3

WBS 1.5.5.1-LANL-007 Systems and Network Equipment

This project ensures the procurement of systems and networking equipment required for support of ongoing production systems. In addition, procurements of related software support and third party labor are coordinated and executed through this project. This project also produces and maintains budgetary planning data and reporting on systems and networking equipment procurement activities.

This project is a coordination point for system and network equipment and related software procurements. Customers include ongoing operations and integration teams, projects, and groups. Interfaces include third party vendors and suppliers and the LANL supply chain and accounting organizations. This project, in coordination with supply chain, negotiates purchase pricing and terms and conditions. Procurements are tracked as they proceed through procurement and fulfillment.

Planned activities in FY08:

- Identify and plan FY08 operational system and network equipment requirements and develop a time phased spend plan, consistent with the budget allocation
- Execute time phased plan; procure and ensure delivery of all required operational system and network equipment including additional Panasas shelves for Roadrunner

Expected deliverables in FY08:

- New networking, storage, and related equipment procurement in support of Roadrunner Phase 3 and TLCC computing systems

Preliminary planned activities in FY09:

- Procuring and ensuring delivery of equipment required in FY09 and planning for all equipment required in FY10 and beyond

WBS 1.5.5.1-LANL-008 Ongoing Facilities

LANL's infrastructure support is handled by three projects: the Nicholas Metropolis Center for Modeling and Simulation (known as the SCC), the Laboratory Data and Communication Center (LDCC), and the Central Computing Facility (CCF). While the majority of facilities work in focused on the SCC and the LDCC buildings there are still a few legacy systems in part of the CCF that need facility support. This effort also covers the cost of electricity for the capacity and capability computing systems, the storage and networking supporting infrastructure, as well as the engineering and facilities work needed to maintain the SCC and LDCC.

Planned activities in FY08:

- Continued operations and maintenance of electrical and mechanical systems for ongoing programmatic computing, providing the facility support for decommissioning the rest of the Q capability system as well as the C capacity systems
- Site preparations for the third phase of additional Roadrunner computer installations.

Expected deliverables in FY08:

- Site Prep for the TLCC systems in the SCC and the LDCC
- Site Prep for Roadrunner Phase 3 system in the SCC
- The 2.4MW SCC Upgrade Project, funded in FY07, will be completed in FY08

Preliminary planned activities in FY09:

- Continued operation of the SCC, LDCC, and CCF facilities and computer rooms
- Enhancements to electricity and cooling for the LDCC to support expected increase in computing capacity along with new system cooling requirements

WBS 1.5.5.1-SNL-001 Production Computing Services

This projects goals are to operate and maintain all production platforms and associated support systems, and operate ASC capability and capacity platforms, data services and visualization systems, long-term hierarchical storage services, high performance network systems, tri-lab compatible cyber authentication and authorization systems, monitoring and reporting services. It will support tri-lab capability platform resource allocations and coordinate with tri-lab peers in establishing priority scheduling if required. Support of Tripod common service and environment decisions and configuration management activities will also be provided.

Planned activities in FY08:

- Deliver production computing service for tri-lab ASC program through continued maintenance and operation of:
  - Red Storm Capability Platform
  - Thunderbird Capacity Platform
  - TLCC system components acquired in FY07 and FY08
  - Red Rose and Black Rose Data Services and Visualization Platforms
  - Long-Term Hierarchical Storage Systems
- Ensure sufficient capacity and capability compute cycles to support ASC Level 2 milestones
- Deploy Tripod capabilities for ASC capacity computing environment
- Emphasis will be made on deploying system-monitoring tools for RAS
- In addition, modify the infrastructure and support environment by:
  - Installing new capacity platforms acquired by the TLCC project and bringing them into production status as rapidly as possible

- Retire end of life IBM Tape robots and StorageTek Powderhorn systems, and begin transition to new Sun/StorageTEK 8500 libraries
- Increase storage capacity of Red Storm disk system

Expected deliverables in FY08:

- Release Kerberos 1.4.X and OpenSSH for distribution; provide Smartcard authentication capabilities to Kerberos services in the unclassified environment
- Deploy NFSV4 server for production use by EAWeb and NWCHub
- Provide tri-labs NFSV4/pNFS services
- Continue build out of Red Storm HPC network environment using equipment purchased in FY07
- Integrate up to 1 petabyte of new disk for Red Storm
- Implement a Tripod common software environment on the TLCC delivered systems
- Complete installation of BlackRose visualization system including increased storage capacity. Replace Visualization nodes on RedRose

Preliminary planned activities in FY09:

- Continued operation of production systems
- Installation and production of TLCC capacity systems

WBS 1.5.5.1-SNL-002 Facilities and Infrastructure

The Facilities and Infrastructure project will:

- Manage all computing facilities and support infrastructure, and provide funds for physical security and utilities (power, cooling, floor space) expenses.
- Plan and coordinate facilities construction or expansion. Design and procure power and cooling equipment as required for production platforms. Supply physical security control for computing facilities and classified media. Assist with tours and audits. Provide 24 hour per day presence in computer center and act as out of hours contact for Scientific Computing.

Planned activities in FY08:

- Develop plan for building expansion project to add 20,000 square feet of additional computer floor space in Building 725, Tech Area I
- Assist in installation and removal of computing equipment as necessary to maintain the Operational capabilities of the ASC environment
- Manage retirement and destruction of classified disk and tape media

Expected deliverables in FY08:

- Power, cooling and facility modifications to support TLCC installations
- Power and cooling for HPSS storage equipment, STK 8500s
- Power and cooling for Red Storm disk system upgrade

Preliminary planned activities in FY09:

- Power, cooling and facility modifications to support future-year TLCC installations

- Begin execution of Building 725 expansion project

WBS 1.5.5.1-SNL-003 Tri-Lab System Integration and Support

The Tri-Lab System Integration and Support project manages projects relating to tri-lab production services and related infrastructure. SNL provides the coordination, operational support and oversight necessary to develop and operate the ASC WAN as an effective resource for the tri-lab community through continued management of the communication link contracts. SNL takes the lead in integrating new encryptor technology into the WAN by obtaining early engineering samples, performing laboratory tests and organizing tri-lab wide functional testing prior to deployment.

Traffic engineering and modeling systems, as well as a dedicated test laboratory based WAN development environment, are used to improve network efficiency and improve utilization. Specially engineering monitoring and management systems are used to analyze network performance and validate vendor provided availability statistics to ensure proper credits are applied to the communication link contracts.

Operating the WAN system overseeing the Qwest communication link contract and monitoring of ESNET connectivity. Provide system level analyst support for cross-site production services relating to data transfer, applications at a distance, access methods and services and coordinate production requests for tri-lab resources. Manage and coordinate SNL participation in the Tripod and TLCC requirements definition and scheduling processes.

Planned activities in FY08:

- Tier 2 and Tier 3 system support personnel will perform problem analysis and problem solution services for issues referred by Customer Support and Production Operations personnel or customers within the tri-lab community.
- Coordinate TLCC and Tripod software activities and ensure configuration management of TLCC.
- Support Production Computing platform monitoring and notification requirements in the Cyber Enterprise Management project.
- Operate ASC WAN loop between LLNL, LANL, and SNL, and manage the Qwest communication link contract ensuring that proper credits are applied to the monthly payments. Coordinate the installation of the SNL designed IPv6 ASC WAN capability in the tri-lab community satisfying the FY08 OMB mandate. Complete the design and deployment of the SNL initiative to provide backup ESnet access for the tri-lab community through the ASC WAN.
- Enhance monitoring of network traffic and cross-site services and access procedures to improve reliability of overall system.

Expected deliverables in FY08

- New WAN Encryptors (10GE (Gigabit Ethernet)) installed in a redundant configuration providing enhanced network availability
- Ipv6-based ASC WAN network transport capability operational
- Emergency tri-lab ESnet ISP access operational via the ASC WAN
- Automated account process integrating SARAPE and WebCARS provides NWC wide support
- Tripod software stack deployed for TLCC07 systems

Preliminary planned activities in FY09:

- Highly integrated WAN and Cyber Enterprise Management monitoring of local and remote production platforms

WBS 1.5.5.2: User Support Services

This level 4 product provides users with a suite of services enabling effective use of ASC tri-lab computing resources. The scope of this product includes planning, development, integration and deployment, continuing product support, and quality and reliability activities collaborations. Projects and technologies include computer center hotline and help-desk services, account management, web-based system documentation, system status information tools, user training, trouble-ticketing systems, and application analyst support.

User Support Services Deliverables for FY08

- TLCC documentation and training modules
- Complete transition to new identity/account management tools
- TLCC training and documentation for Alliance users and other ASC users
- Roadrunner Phase 3 training in preparation for hybrid computing
- Deliver reliable and responsive service to users in the ASC tri-lab computing environments

WBS 1.5.5.2-LLNL-001 Hotlines and System Support

The Hotlines and System Support project provides users with a suite of services enabling effective use of ASC tri-lab computing resources.

This project includes computer center hotline and help-desk services, account management, web-based system documentation, system status information tools, user training, trouble-ticketing systems, and application analyst support. Services are supplied to users from external sites including LANL, SNL, and the ASC Alliance sites, as well as the LLNL users.

Planned activities in FY08:

- Design and implementation of new identity/account management tools
- Development of TLCC documentation and training
- Special training for new Alliance partners
- Investigation of new institutional trouble ticket, with emphasis on ASC requirements
- Ongoing support services for hotline operations, documentation, and training

Expected deliverables in FY08:

- TLCC documentation and training modules
- Complete transition to new identity/account management tools

Preliminary planned activities in FY09:

- Development of Sequoia ID system documentation and training
- Ongoing support services for hotline operations, documentation, and training

WBS 1.5.5.2-LANL-001 Integrated Computing Network Consulting, Training and Documentation, and External Computing Support

This project is comprised of three teams:

- The Integrated Computing Network consulting office, responsible for direct customer service to local and remote users of LANL ASC resources
- The training and documentation team, responsible for the development and delivery of documentation and training materials for LANL ASC resources
- The external computing support team, responsible for usage statistics and the administrative interface for external HPC users such as tri-lab ASC and ASC Alliances

Thrust areas consist of user support services, operational metrics for an HPC environment on, for example, usage and availability, Web page development to present this information to system personnel and users, and the development of user documentation and training.

Planned activities in FY08:

- Perform ongoing user support for users of LANL's ASC computing resources
- Expand online documentation in both breadth and depth
- Deploy new trouble tracking systems in the open and secure
- Deploy enhanced HPC statistical reports and infrastructure
- Re-architect the accounts system
- Further refinement and development of ongoing customer satisfaction surveys to quickly identify and deal with any future shortcomings

Expected deliverables in FY08:

- New training classes and documentation developed for users of TLCC clusters with Tripod stack
- New training classes for Roadrunner computing with focus on future hybrid computing

Preliminary planned activities in FY09:

- Perform ongoing user support for users of LANL's ASC computing resources
- Continue improving the quality of support for productive use of the computing services

WBS 1.5.5.2-LANL-002 Ongoing User Services

This project has three main focus areas grouped under the umbrella of user services. The first focus is to support the production environments by ensuring that tools needed by users are available and working on production platforms. Tools in this context include compilers, debugging tools, and performance and profiling tools. The second focus is on application development and end user in-person support. This involves technical consulting and collaborations with customers on issues related to the application code development environment and use on production systems, including documentation and training. The third focus is on application readiness of systems. This work is coordinated under FOUS and CSSE to support the applications readiness team at LANL that works directly with users to address application problems with the ASC computers.

This team supports users/code teams to do initial ports and deep problem determination of code issues on new platforms. The intent of the team is to supply computer scientists that know both applications and systems.

Planned activities in FY08:

- Provide support for user software depending on the product and needed support. At the low end, products will be installed and tested to ensure they are operational on platforms. At the higher end, development work may be included or coordinated to make the products work in our environment and to enhance for better utilization.
- Continued process improvement with HPC-3 and HPC-4 teams for product installation, license server management, user issue, and support. The target is to get effective communication systems in place and partnerships for support.
- Software packages will be managed through a change control repository; RPMs will be made for installation where possible.
- A nightly (or periodic) regression test system for tools will continue to be implemented. This is to ensure the health of the software tools environment and track performance points.
- Continue improving a monitoring and tracking system that checks and compares tool versions running on platforms and segments. In addition, usage statistics will be gathered for some tools.
- Third-party requests will be reviewed and supported once vetted with appropriate approval bodies. Some support will be provided to teams that want to install specific local tools.
- Ensure an acceptance/regression test suite process is initiated for tool acceptance prior to production implementation.
- Provide technical consulting and collaborations with customers on issues related to the application code development environment on production systems.
- Manage requests for and installations of third-party software on the production systems.
- Focus on the applications targeted for the Roadrunner Capacity System.

Expected deliverables in FY08:

- Ongoing user services will continue to be delivered in FY08. No specific new deliverables are planned in FY08. However, work will start on developing tools and services that will be specific deliverables in FY09 to support user services for Roadrunner Phase 3 hybrid computing.

Preliminary planned activities in FY09:

- Planning and developing for software tools required for future systems especially Roadrunner Phase 3.
- Continuing to improve the quality of support for productive use of the computing services. This includes continuing the role of the application readiness team to help ensure issues with optimal use of ASC compute resources are being addressed.

WBS 1.5.5.2-SNL-001 User Environment and Application Support

The user environment and application support teams facilitate computing on ASC tri-lab platforms, as well as SNL computing systems.

The user environment provides information, tools, training resources, and direct user support pertaining to the use of the ASC scientific computing environment. Information resources include access to a knowledge management and retrieval system through the Collaborative Learning, Information, and Knowledge web-enabled tool. User support is provided through a combination of prime-time telephone support and email support. Activities related to tri-lab use include: assistance to tri-lab customers with problems or requests at SNL; assistance to SNL customers with their computing at remote locations; participation in the management of SNL computing resources via the CCC process for Purple and the Sandia Platform Oversight Committee process at SNL; and representing SNL needs to the EPR process.

The application support team works with a breadth of ASC applications and system environments to develop and apply expertise that enables efficient and effective use of ASC's precious computing resources. The team provides porting assistance to new ASC architectures; benchmarking; and various forms of application performance analysis, including modeling, scaling studies, and optimization.

Planned activities in FY08:

- Provide reliable and responsive service to the users in the ASC tri-lab and SNL computing environments
- Perform scaling studies of SNL's critical ASC applications on Red Storm and Purple
- Coordination of SNL computing requirements via tri-lab EPR meetings, CCC process for Purple, and Sandia Platform Oversight Committee process at SNL

Expected deliverables in FY08:

- Deliver reliable and responsive service to users in the ASC tri-lab and SNL computing environments
- Deliver improved training for system administrators and users via web-based short courses

Preliminary planned activities in FY09:

- Improve reliability and responsiveness in service to users in the ASC tri-lab and SNL computing environments
- Use system models to assist in evaluation of future advanced architecture options with respect to performance of SNL applications
- Continue refinement of code performance on Red Storm Dual Core system and other ASC platforms, such as Purple
- Coordination of SNL computing requirements via tri-lab EPR meetings, CCC process for Purple, and Sandia Platform Oversight Committee process at SNL

WBS 1.5.5.3: Collaborations

This level 4 product provides collaboration with external agencies on specific high-performance computing projects. The scope of this product includes planning, development, integration and deployment, continuing product support, and quality and reliability activities collaborations. This product also includes any programmatic support across the entire ASC program and studies, either by internal or external groups, that enable the program to improve its planning and execution of its mission.

Collaborations Deliverables for FY08

- Provide all necessary support for the SC07 ASC research exhibit conference booth
- Conceptualize, develop, provide, and manage all multi-media and written communications for NA-114
- Manage the ASC Russian Program
- Organize and support the annual tri-lab Principal Investigators (PI) meeting

WBS 1.5.5.3-LLNL-001 Program Support

The Program Support project provides service to the ASC program. These services include procurement and contracting, project management, and meeting support. These services are in support of both LLNL-only and tri-lab activities.

Planned activities in FY08:

- Procurement activities will include management of existing tri-lab contacts and negotiating and executing new contracts
- Meeting support will be provided for the ASC PI meeting, the annual Supercomputing conference, Predictive Science Panel meetings, and other meetings and workshops

Expected deliverables in FY08:

- Continued management of existing Purple and BlueGene/L contracts
- Execution and management of the TLCC07 contract
- Negotiation and execution of Sequoia contract

Preliminary planned activities in FY09:

- Continuation of the FY08 activities

WBS 1.5.5.3-LANL-001 Program Support

Through the Program Support project, LANL provides support to the national program, both by providing resources and expertise to the federal program office and by participating in coordination and integration activities for the tri-lab program.

Planned activities in FY08:

- Alternate with Livermore in hosting the Predictive Science Panel. Results will be incorporated into Program plans and initiatives. The Predictive Science Panel consists of outstanding academic experts in the field of predictive science, and helps to build relations with the academic and industrial community of interest to predictive science.
- Provide consultant support to the federal program management efforts to foster collaborations and build support within the predictive science community.
- Work with the lead lab in FY08 in organizing the annual tri-lab PI meeting.

Expected deliverables in FY08:

- Support for SC07 conference
- Support for ASC PI meeting

Preliminary planned activities in FY09:

- LANL will host the Predictive Science Panel in FY09

WBS 1.5.5.3-SNL-001 One Program/Three Labs

One Program/Three Labs funds several critical coordination and integration activities essential to the success of ASC. These are divided into four distinct parts: 1) provide ASC multi-level communications per existent communications plan and by special request; 2) support for tri-lab and Headquarters interactions including ASC executive meetings, the ASC principal investigators meetings, and an SAIC contract for Headquarters administration support; 3) SNL outreach to the DoD laboratories and programs; and 4) collaborative projects in modeling and simulations with the Russian nuclear laboratories and institutes. These crosscut and outreach activities seek to facilitate cooperation and collaboration amongst U.S. and international laboratories, improve program visibility within the high-performance computing community and enhance the overall operations of the ASC program.

Planned activities in FY08:

- Continued development of the SNL and Headquarters ASC website
- Continued production of high-quality communications materials for Headquarters and the broader high-performance computing community
- Special assignments for Complex 2030 such as the S&T roadmap (as required) and special events (expositions)
- Support for the ASC Executive Committee; support for quarterly meetings of the ASC executive committee; support for the annual PI meetings that expose attendees to technical and programmatic efforts at the three laboratories; management of the SAIC contract to provide various administration support to Headquarters
- Project review meeting with the Russian Federal Nuclear Laboratories and various National Academy Institutes
- Act as lead in FY08 for organizing the annual tri-lab PI meeting; FY08 activities will continue to be a tri-lab program management effort.

Expected deliverables in FY08:

- Conceptualize, develop, provide, and manage all multi-media and written communications for NA-114
- Manage the ASC Russian Program
- Continue to support cooperative work with the SC07 ASC Research Exhibit
- Act as lead in FY08 for organizing the annual tri-lab PI meeting

V. ASC Level 1 and 2 Milestones

Table V-1. Quick Look: *Proposed* Level 1 Milestone Dependencies

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)	ASC Category	Depends on another Milestone	Milestone ID	Has another Milestone depending on it	Milestone ID
359	Complete modern baseline of all enduring stockpile systems with ASC codes.	1	2009	Sep-09	ASC	HQ, LLNL, LANL, SNL	DSW Deliverables		349		
1	Develop, implement, and apply a suite of physics-based models and high-fidelity databases to enable predictive simulation of the initial conditions for secondary performance.	1	FY09	Q4	ASC	HQ, LLNL, LANL, SNL	C11, C4				
2	Develop, implement, and validate a suite of physics-based models and high-fidelity databases in support of Full Operational Capability in DTRA's National Technical Nuclear Forensics program.	1	FY09	Q4	ASC	HQ, LLNL, LANL	C11, C1, C4, NA-22, DTRA				
3	Baseline demonstration of UQ aggregation methodology for full-system weapon performance prediction	1	FY10	Q4	ASC	HQ, LLNL, LANL, SNL	C11, C1, C4, DSW				
4	Develop, implement, and apply a suite of physics-based models and high-fidelity databases to enable predictive simulation of the initial conditions for primary boost.	1	FY12	Q4	ASC	HQ, LLNL, LANL	C11, C1, C2				

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)	ASC Category	Depends on another Milestone	Milestone ID	Has another Milestone depending on it	Milestone ID
5	Capabilities for significant finding investigation response improvements	1	FY13	Q4	ASC	HQ, LLNL, LANL, SNL	C11, DSW				
6	Develop, implement, and apply a suite of physics-based models and high-fidelity databases to enable predictive simulation of primary boost	1	FY15	Q4	ASC	HQ, LLNL, LANL	C11, C1, C2, C10				
7	Develop predictive capability for full-system integrated weapon safety assessment	1	FY16	Q4	ASC	HQ, LLNL, LANL, SNL	C11, C1, C2, DSW				
8	Develop, implement, and apply a suite of physics-based models and high-fidelity databases to enable predictive simulation of secondary performance	1	FY20	Q4	ASC	HQ, LLNL, LANL, SNL	C11, C4, C2, C10				

Table V-2. Quick Look: Level 2 Milestone Dependencies for FY08[6]

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)	ASC Category	Depends on another Milestone	Milestone ID	Has another Milestone depending on it	Milestone ID
	Deliver a physics-based sub-grid model to support energy balance resolution	2	08	Mar-08	ASC/Integrated Codes	LLNL					
	Demonstrate progress toward ALE/AMR hydrodynamic capability	2	08	Mar-08	ASC/Integrated Codes	LLNL					
	Evaluate opportunities for sharing of modules across ASC codes	2	08	Jun-08	ASC/Integrated Codes	LLNL					
	Assess adequacy of current models for the initial conditions for boost	2	08	Sep-08	ASC/Integrated Codes	LLNL					
	Explore and assess opportunities to enhance the nuclear weapon simulation capability of performance code for future stockpile activities	2	08	Sep-08	ASC/Integrated Codes	LLNL					
	Develop enhanced capabilities in a multi-dimensional (1D, 2D, and 3D) effects code to support current and planned effects modeling efforts	2	08	Sep-08	ASC/Integrated Codes	LLNL					
	Improve physics models and nuclear data to address known deficiencies and facilitate the expansion of our nuclear forensic database to include emplacement scenarios	2	08	Sep-08	ASC/Integrated Codes	LLNL					

[6] Factors such as FY08 Congressional Appropriations, NNSA/DP directives, and National Security considerations may necessitate a change in the current milestone set.

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)	ASC Category	Depends on another Milestone	Milestone ID	Has another Milestone depending on it	Milestone ID
	Enhance the capabilities and extend the advanced multi-scale material strength model to another relevant material	2	08	Sep-08	ASC/Physics and Engineering Models	LLNL					
	New baseline global EOS data library delivered for QMU, V&V, and other applications	2	08	Sep-08	ASC/Physics and Engineering Models	LLNL					
	Advanced EOS tables for NIF ICF capsule materials delivered for capsule design, QMU, and V&V	2	08	Sep-08	ASC/Physics and Engineering Models	LLNL					
	Assessment of nuclear physics uncertainties	2	08	Sep-08	ASC/Physics and Engineering Models	LLNL					
	Use of the UQ methodology to provide predictive capability in off-normal situations	2	08	Sep-08	ASC/Verification & Validation	LLNL					
	Validation of energy balance model	2	08	Sep-08	ASC/Verification & Validation	LLNL					
	Contractor software quality assurance audit of ASC codes	2	08	Sep-08	ASC/Verification & Validation	LLNL					
	Estimated ranges of adjustable parameters in an ASC code (40%)	2	08	Sep-08	ASC/Verification & Validation	LLNL					
	PMP set expanded to 25 events	2	08	Sep-08	ASC/Verification & Validation	LLNL					
	Deploy Moab resource management services on BlueGene/L	2	08	Jun-08	Computational Systems & Software Environment, Facility Operations & User Support	LLNL					
	Deploy identity/account management automation tools	2	08	Sep-08	Computational Systems & Software Environment, Facility Operations & User Support	LLNL					

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)	ASC Category	Depends on another Milestone	Milestone ID	Has another Milestone depending on it	Milestone ID
	Update and document Software Life Cycle Plan for the LANL ASC modern multi-physics code projects	2	08	Dec-07	ASC/Integrated Codes	LANL					
	Release of a Crestone project code to support X Program requirements, attribution, TBI and Campaign 4	2	08	Mar-08	ASC/Integrated Codes	LANL					
	Release of a Shavano project code to support X Program requirements, attribution, TBI and Campaign 4	2	08	Mar-08	ASC/Integrated Codes	LANL					
	High-fidelity simulation of a lighting system using a Crestone project code	2	08	Sep-08	ASC/Integrated Codes	LANL					
	Release of a Crestone project code to support high fidelity simulations of a lighting system	2	08	Sep-08	ASC/Integrated Codes	LANL					
	Release of a Shavano project code to support X Program requirements, attribution, TBI and dynamic plutonium experiments	2	08	Sep-08	ASC/Integrated Codes	LANL					
	Enhanced Pu multiphase EOS capability	2	08	Sep-08	ASC/Physics & Engineering Models	LANL					
	Develop, benchmark, and compare stand-alone NLTE opacity capabilities	2	08	Sep-08	ASC/Physics & Engineering Models	LANL					
	Fundamental simulations of material response (MD code) and plasma physics (PIC code) on HPC platforms in support of resolving one of the "knobs"	2	08	Sep-08	ASC/Physics & Engineering Models	LANL					

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)	ASC Category	Depends on another Milestone	Milestone ID	Has another Milestone depending on it	Milestone ID
	ASC simulations supporting the National technical Nuclear Forensics attribution program	2	08	Sep-08	ASC/Physics & Engineering Models	LANL					
	Physics enhancements to ductile damage/failure framework to account for 3D expansion deformation	2	08	Sep-08	ASC/Physics & Engineering Models	LANL					
	V&V assessment of ASC codes for thermonuclear applications, including initial definition of secondary validation test suite	2	08	Dec-07	ASC/Verification & Validation	LANL					
	V&V assessment of late-time primary implosion	2	08	Sep-08	ASC/Verification & Validation	LANL					
	Complete an assessment of the Roadrunner final system	2	08	Dec-07	ASC/Computational Systems & Software Environment	LANL					
	Take delivery of Roadrunner Phase 3 system	2	08	Sep-08	ASC/Facility Operations & User Support	LANL					
	Predictive failure capabilities in SIERRA mechanics	2	08	Sep-08	ASC/Integrated Codes	SNL					
	Neutron tube source plasma generation and transport	2	08	Sep-08	ASC/Integrated Codes	SNL					
	Demonstrate capability to predict combined cable SGEMP, box SGEMP, and cavity SGEMP events using the RAMSES modules Ceptre, ITS, and Emphasis for WR1.	2	08	Q1	ASC/Integrated Codes	SNL					

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)	ASC Category	Depends on another Milestone	Milestone ID	Has another Milestone depending on it	Milestone ID
	Develop high-fidelity model of the SGT/AT accident analysis; and using the Presto code, demonstrate simulation capability (including visualization) at scale on Red Storm/Purple capability computing platform.	2	08	Q1	ASC/Integrated Codes	SNL					
	Deploy implicit contact enforcement in ACME	2	08	Sep-08	ASC/Integrated Codes	SNL					
	Demonstration of Geometry Tolerant Mesh Generation	2	08	Sep-08	ASC/Integrated Codes	SNL					
	Combined effects (age aware) temperature/radiation model of complementary metal oxide semiconductor (CMOS) integrated circuit technology.	2	08	Sep-08	ASC/Physics & Engineering Models	SNL					
	Initial model demonstration for general component foam encapsulation processing	2	08	Sep-08	ASC/Physics & Engineering Models	SNL					
	Validation of thermal models for a prototypical MEMS thermal actuator	2	08	Sep-08	ASC/Physics & Engineering Models	SNL					
	Uncertainty analysis for an RB electrical cable in radiation environment	2	08	Sep-08	ASC/Verification & Validation	SNL					
	Deliver post-processing tools that enable verification and validation of FY08 simulations	2	08	Sep-08	ASC/Computational Systems & Software Environment	SNL					
	Infrastructure deployment plan for ASC petascale environments	2	08	Mar-08	ASC/Computational Systems & Software Environment, Facility Operations & User Support	LLNL, LANL, SNL					

Table V-3. Quick Look: *Preliminary* Level 2 Milestone Dependencies for FY09

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)	ASC Category	Depends on another Milestone	Milestone ID	Has another Milestone depending on it	Milestone ID
	Deliver improved AGEX capabilities to users	2	09		ASC/Integrated Codes	LLNL					
	Improve physics models and nuclear databases in support of nuclear forensic full operational capability	2	09		ASC/Integrated Codes	LLNL					
	Explore and assess required improvements in numerical treatments of materials failure and fracture	2	09		ASC/Integrated Codes	LLNL					
	Advanced model of HE initiation and failure	2	09		ASC/Physics & Engineering Models	LLNL					
	Advanced radiative properties model	2	09		ASC/Physics & Engineering Models	LLNL					
	Next generation hydrodynamic model	2	09		ASC/Physics & Engineering Models	LLNL					
	Multi-scale material strength model with applications	2	09		ASC/Physics & Engineering Models	LLNL					
	Standard calculation 2009 (SC09)	2	09		ASC/Verification & Validation	LLNL					
	Estimated ranges of adjustable parameters in an ASC code (80%)	2	09		ASC/Verification & Validation	LLNL					
	Production Deployment of HPSS Release 7.1	2	09		ASC/Computational Systems & Software Environment	LLNL					
	Application code correctness tool suite for 10,000 or more processors	2	09		ASC/Computational Systems & Software Environment	LLNL					

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)	ASC Category	Depends on another Milestone	Milestone ID	Has another Milestone depending on it	Milestone ID
	Sequoia Phase 1 (Dawn ID) I/O and archival infrastructure initial deployment	2	09		ASC/Computational Systems & Software Environment	LLNL					
	Preliminary constitutive framework for representation of certain material	2	09		ASC/Physics & Material Models	LANL					
	Documentation of current mix models and evaluation against a validation test suite	2	09		ASC/Physics & Material Models	LANL					
	Demonstrate a scientific application that uses a significant portion of the hybrid Roadrunner system	2	09		ASC/Physics & Material Models	LANL					
	National Technical Nuclear Forensics with DTRA involving Shavano Project coupled with isotopics coupled with neutronics	2	09		ASC/Physics & Material Models	LANL					
	Validation assessment of penetration mechanics for surety applications	2	09		ASC/Verification & Validation	LANL					
	Catalog of Major Adjustable Parameters in Weapons Physics Simulations	2	09		ASC/Verification & Validation	LANL					
	Performance assessment of the full Roadrunner system	2	09		Computational Systems & Software Environment	LANL					
	Coupled thermal structural capability in SIERRA mechanics to model structural collapse during a fire	2	09		ASC/Integrated Codes	SNL					
	Initial SIERRA mechanics capability of coupled PREMO/CALORE ablation capability for reentry environments	2	09		ASC/Integrated Codes	SNL					

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)	ASC Category	Depends on another Milestone	Milestone ID	Has another Milestone depending on it	Milestone ID
	Improved energy dissipation models for predictive mechanical response	2	09		ASC/Integrated Codes	SNL					
	Algorithms for system-level uncertainty assessment	2	09		ASC/Integrated Codes	SNL					
	Demonstrate capability to build and manage massive meshes	2	09		ASC/Integrated Codes	SNL					
	Improved pressure fluctuation model for turbulent flow	2	09		ASC/Physics & Material Models	SNL					
	Initial validation of weapon safety models for abnormal mechanical environments	2	09		ASC/Verification & Validation	SNL					
	Evaluate advanced memory subsystems	2	09		ASC/Computational Systems & Software Environment	SNL					
	Deliver smart simulation/experimental petascale data comparison tools	2	09		ASC/Computational Systems & Software Environment	SNL					
	Evaluation of the impact chip multiprocessors have on Sandia application performance	2	09		ASC/Computational Systems & Software Environment	SNL					
	Deploy Tripod capabilities for capacity computing environment	2	09		ASC/Facility Operations & User Support	LLNL, LANL, SNL					

Detailed Milestone Descriptions for FY08

Milestone (ID#): Deploy Moab resource management services on BlueGene/L
Level: 2
Fiscal Year: FY08
DOE Area/Campaign: ASC
Completion Date: Jun-08
ASC nWBS Subprogram: Computational Systems & Software Environment, Facility Operations & User Support
Participating Sites: LLNL
Participating Programs/Campaigns: ASC
Description: In September 2006, the Moab Workload Manager was selected to become the standard batch scheduling system for exclusive use across the tri-lab HPC facilities. Moab is a commercial product that is developed and sold by Cluster Resources, Inc. LLNL's existing batch system, LCRM, will gradually be replaced by Moab on all platforms. Moab workload management services were installed on several LLNL platforms in early FY07. The BlueGene/L system currently runs the SLURM resource manager and is scheduled by LCRM. This milestone represents LLNL efforts to enhance both SLURM and Moab to extend Moab's capabilities to schedule and manage BlueGene/L, and increases portability of user scripts between ASC systems.
Completion Criteria: This milestone is complete when batch jobs can be specified, submitted to Moab, scheduled and run on the BlueGene/L system. Moab will be able to support the markedly increased scale in node count as well as the wiring geometry that is unique to BlueGene/L. Moab will also prepare and report statistics of job CPU usage just as it does for the current systems it supports.
Customer: NNSA/ASC Headquarters
Milestone Certification Method: Completion evidence for this milestone will be in the form of (1) documentation - a report that certifies that the completion criteria have been met; and (2) user hand-off.
Supporting Resources: Tri-lab CSSE and FOUS products and personnel with the support from CRI to add the necessary enhancements to Moab.
Codes/Simulation Tools Employed: TBD
Contribution to the ASC Program: Enhances effectiveness of the infrastructure to run ASC codes on the BlueGene/L platform(s).
Contribution to Stockpile Stewardship: Supports the overall SSP goal that rely on codes that run on BlueGene/L, including UQ analyses, advanced weapons science studies, and enhanced predictive capabilities.

No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1.	This milestone has dependencies on an external vendor (Cluster Resources) to provide topology-aware capabilities and scalability for large BlueGene/L systems.	Low	Moderate	Low

Milestone (ID#): Deploy identity/account management automation tools
Level: 2
Fiscal Year: FY08
DOE Area/Campaign: ASC
Completion Date: Sep-08
ASC nWBS Subprogram: Computational Systems & Software Environment, Facility Operations & User Support
Participating Sites: LLNL
Participating Programs/Campaigns: ASC
Description: The LC LLNL HPC Identity and Account management system will replace the paper system used by users to request accounts, and replace current backend systems (IAM and LCAMS), currently used to process the requests, with a Web-based system. This deliverable will produce a flexible Identity Management system that can streamline the account request, approval and instantiation process, while simultaneously improving security, and reducing costs associated with this activity. This milestone deliverable will interface with the LLNL Institutional Identity Management system to leverage data held by them relating to physical site access, employment status and remote access. While, some components are expected to be deployed in Q1 and Q2 full functionality for a majority of LC users should be attained in Q3 to complete the milestone.
Completion Criteria: The IDM system is generally available and has been used for the creation of accounts on at least two ASC resources at LLNL. Used by at least one LLNL user who is an LLNL employee, and one collaborator (ASC inter-site user) utilizing the web interface for both the request and approvals, as well as automated provisioning. Demonstration of audit process to the LLNL Computer Information System Security Officer, and acceptance of the audit trail by the LLNL Computing ISSO.
Customer: All users of ASC computer systems hosted by LLNL.
Milestone Certification Method: Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. The "handoff" of the developed capability (product) to LC management.
Supporting Resources: Computational Systems & Software Environment, Facility Operations & User Support
Codes/Simulation Tools Employed: N/A
Contribution to the ASC Program: Efficiency and security by modernizing management of accounts and resources by eliminating paper and wet signatures with an all electronic process.
Contribution to Stockpile Stewardship: Promotes efficiency and security by modernizing management of accounts and resources.

No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1.	LSO has less history with electronic records than with paper records, which may pose approval and process hurdles.	Low	Moderate	Low
2.	Potential issues with the quality of data provided by both the LLNL Institutional Identity management process and the ICSI/ESN project.	Low	Moderate	Low

Milestone (ID#): Complete an assessment of the Roadrunner final system
Level: 2
Fiscal Year: FY08
DOE Area/Campaign: ASC
Completion Date: Dec-07
ASC nWBS Subprogram: Computational Systems & Software Environment
Participating Sites: LANL
Participating Programs/Campaigns: ASC
Description: The Roadrunner Phase 3 procurement is a contract option. The option is a method of managing risk for HPC systems. As part of the decision to procure the Phase 3 option, a list of 5 performance criteria will be used to evaluate the Phase 3 system. These are the five criteria: I. Will the proposed hybrid system be applicable to the future weapons workload? II. Will the proposed system achieve a sustained petaFLOPS on LINPACK? III. Will the proposed programming model and provided Application Programming Interface (API) provide a reasonable vehicle for programming hybrid systems? IV. Will the system management and integration plan be effective for running the final system at operational scale? V. Is appropriate progress being made on development of IBM technology to be delivered in final system? A review committee consisting of tri-lab, academic and industrial members of the HPC community will review the technical results against the above criteria.
Completion Criteria: This milestone is complete after review is held and summary report is written.
Customer: LANL Management; Roadrunner Executive Team; Roadrunner Federal Project Director; Roadrunner Acquisition Executive
Milestone Certification Method: A program review is conducted and its results are documented. Viewgraphs from presenters, as well as viewgraphs from Review committee. Review committee will also generate report with summary.
Supporting Resources: IC, CSSE
Codes/Simulation Tools Employed: A subset of the following scientific simulation applications will be used for the assessment: Sweep3D, VPIC, Milagro, SPASM, Pagosa
Contribution to the ASC Program: The successful completion of this L2 Milestone would lead to the procurement of the first large hybrid HPC system. This is the future direction of HPC and this would provide a significant resource under which to contribute to ASC leadership in hybrid computing tools and applications.

Contribution to Stockpile Stewardship: This L2 Milestone is an important part of the process in determining to execute the procurement of the Roadrunner final system. This system would provide a significant increase in computing cycles at LANL for SSP applications.				
No.	**Risk Description**	**Risk Assessment (low, medium, high)**		
		Consequence	**Likelihood**	**Exposure**
1	Unable to complete assessment	High	Low	High
2	Lack of Personnel availability	High	Low	High
3	Vendor unable to meet scope or schedule for HW/SW deliverables	High	Medium	Medium
4	Insufficient budget	High	Low	High

Milestone (ID#): Take delivery of Roadrunner Phase 3 system
Level: 2
Fiscal Year: FY08
DOE Area/Campaign: ASC
Completion Date: Sep-08
ASC nWBS Subprogram: Facility Operations & User Support
Participating Sites: LANL
Participating Programs/Campaigns: ASC
Description: The current Roadrunner contract calls for a decision to be made in October 2007, on whether to exercise the option to procure the final hybrid system from IBM. If the decision is made to proceed then the focus of this effort will be to provide all the project management needed to deliver the Phase 3 system to LANL in FY08. The Roadrunner final system is scheduled to deliver a significantly advanced architecture system that should provide compute power of over a petaFLOPS of computing cycles to the weapons program. The advanced architecture hardware will consist of a hybrid computing architecture that has the potential for significant improvements to the price/performance curve to help meet the computing requirements in the future. The goal of this milestone is to have the Roadrunner final system delivered to LANL in preparation for system installation and testing in the fall.
Completion Criteria: Machine has been assembled at the vendor's site and has demonstrated sustained petaFLOPS computing, and the number of CUs as specified in the statement of work have been delivered to LANL.
Customer: Stockpile Stewardship Program, to include weapons science customers
Milestone Certification Method: A program review will be conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, will be prepared as a record of milestone completion.
Supporting Resources: Approximately 14 FTEs
Codes/Simulation Tools Employed: As specified in the statement of work.
Contribution to the ASC Program: Brings compute power for the ASC program into the petaFLOPS range, enabling scientific calculations in support of the ASC Roadmap.
Contribution to Stockpile Stewardship: Provides crucial compute cycles and speed to weapon's simulation and physics codes for taking the predictive capability to the next level.

No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1	Vendor late in delivering system on time	High	Medium	Medium
2	SCC electrical upgrade not complete	Medium	Low	Medium

Milestone (ID#): Deliver post-processing tools that enable verification and validation of FY08 simulations
Level: 2
Fiscal Year: FY08
DOE Area/Campaign: ASC
Completion Date: Sep-08
ASC nWBS Subprogram: Computational Systems & Software Environment
Participating Sites: SNL
Participating Programs/Campaigns: ASC
Description: Provide scalable analysis capabilities to support the V&V of large data with particular emphasis on comparison of large datasets and validation. The project expects to deliver capability supporting specific aspects of large data analysis required by such customers as HEDP (specific capabilities are still under negotiation).
Completion Criteria: Complete a set of V&V analyses that (1) deliver useful results in support of customer applications; and (2) are applicable to a range of customers and problem areas.
Customer: V&V, Integrated Codes
Milestone Certification Method: Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. The "handoff" of the developed capability (product) to a nuclear weapons stockpile customer is documented.
Supporting Resources: Post-Processing development and support staff; V&V and application partner resources to guide feature development and generate application data; computing resources on ASC platforms and/or data analysis servers as needed.
Codes/Simulation Tools Employed: ParaView
Contribution to the ASC Program: Provide important tools needed to enable V&V/QMU for ASC simulations in support of predictive capabilities.
Contribution to Stockpile Stewardship: Provide important tools needed to support simulation-based weapon qualification/certification.

No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1	Inadequate access to V&V and customer partner resources.	High	Low	Medium
2	Inadequate access to simulation result data for V&V analysis; inadequate access to compute resources needed to do V&V simulation runs to produce results.	High	Low	Medium

Milestone (ID#): Infrastructure deployment plan for ASC petascale environments
Level: 2
Fiscal Year: FY08
DOE Area/Campaign: ASC
Completion Date: Mar-08
ASC nWBS Subprogram: Computational Systems and Software Environment (CSSE) & Facility Operations and User Support (FOUS)
Participating Sites: LLNL, LANL, SNL
Participating Programs/Campaigns: ASC
Description: The ASC Petascale Environment Infrastructure Deployment Plan will identify, assess, and specify the development and deployment approaches for critical components in four different technical areas: (1) development environment and tools; (2) petascale data analysis; (3) I/O, file systems and archives; and (4) networks and interconnects. This Plan will identify and quantify potential technical gaps or issues, and, where they exist, will define a prioritized approach to closing those gaps. While the specific deliverable (a planning document) for this milestone is to be completed in Q2 FY08, the petascale infrastructure components will likely be deployed throughout a five-year FY08–FY12 timetable, and prioritization given to enabling predictive weapons simulations. The plan will be applicable to multiple ASC petascale platforms deployed during that time, including Roadrunner and the Sequoia ID and final systems.
Completion Criteria: This milestone will produce an integrated tri-lab planning document that can be applied to multiple ASC petascale environments over multiple years. While initial petascale infrastructures and capabilities may be site or platform specific, the eventual goal as supported by this milestone is to provide a seamless petascale user environment for capability computing as envisioned and motivated by the ASC Roadmap. The deployment plan provided by this milestone will operate as a stand-alone document, but it may be desirable to also incorporate, by reference or as appendices, additional Tri-lab or lab-specific documents, for example the "ASC I/O and Storage FY06-FY10 Technology Update and Plan for HPC File Systems, Scalable I/O, and Archival Storage," and the "LLNL FY08 I/O Integration Blueprint."
Customers: Tri-lab and individual laboratory system integration teams, future users of ASC petascale computing platforms, NNSA/ASC Headquarters and ASC Execs.
Milestone Certification Method: Completion evidence for this milestone will be: (1) professional documentation - the formal plan, reviewed and released for wide (or potentially unlimited) distribution; and (2) an appropriate internal program review with documented results, incorporating any supporting presentation slides.
Supporting Resources: Tri-lab CSSE and FOUS products and project team personnel.
Codes/Simulation Tools Employed: TBD
Contribution to the ASC Program: Provides the detailed deployment planning for petascale computational environments and infrastructure, targeted at improving the usability and cost performance of petascale capability platforms and codes.

Contribution to Stockpile Stewardship: Ensures successful deployment of petascale computational environments in support of overall SSP goals, including Uncertainty Quantification (UQ) analyses, advanced weapons science studies, and enhanced integrated design code predictive capabilities.				
No.	**Risk Description**	**Risk Assessment (low, medium, high)**		
		Consequence	**Likelihood**	**Exposure**
1.	Technical working groups may not reach consensus or meet planned schedules for first and final drafts of document. Risk control through Lab POCs and HQ coordination.	Low	Moderate	Low
2.	Some technical gaps may require complex multi-lab and/or multi-agency collaborative approaches, although this will likely not delay creating the planning document. Risk control through Lab POCs and HQ coordination.	Low	Moderate	Low
3.	Inability to predict the future state of some petascale technologies may cause document to not be as effective as envisioned. Technical working groups, POCs and HQ to accept this risk.	Low	Moderate	Low

Milestone Descriptions for Preliminary FY09

Milestone (ID#): Production Deployment of HPSS Release 7.1
Level: 2
Fiscal Year: FY09
DOE Area/Campaign: ASC
Completion Date: Q4
ASC nWBS Subprogram: Computational Systems & Software Environment
Participating Sites: LLNL
Participating Programs/Campaigns: ASC
Description: HPSS Release 7.1, which is sharply focused on file create and transactional performance in support of petascale environments, will be deployed on both unclassified and classified production networks through two carefully planned rollout efforts: (1) OCF production by the end of Q2FY09; and (2) SCF production by the end of Q4FY09.

Milestone (ID#): Application code correctness tool suite for 10,000 or more processors
Level: 2
Fiscal Year: FY09
DOE Area/Campaign: ASC
Completion Date: Q3
ASC nWBS Subprogram: Computational Systems & Software Environment
Participating Sites: LLNL
Participating Programs/Campaigns: ASC
Description: A code correctness tools suite will be deployed that allows users to debug a variety of issues that arise when running at scale. The suite will include tools that, when applied to jobs running on 10,000 plus processors, quickly identify information that narrows the problem space and supports applying other tools in the suite to jobs at scale. Tool suite capability will be demonstrated on the Sequoia Project's initial delivery and code scalability platform (Dawn ID).

Milestone (ID#): Sequoia Phase 1 (Dawn ID) I/O and archival infrastructure initial deployment
Level: 2
Fiscal Year: FY09
DOE Area/Campaign: ASC
Completion Date: Q1
ASC nWBS Subprogram: Computational Systems & Software Environment
Participating Sites: LLNL
Participating Programs/Campaigns: ASC
Description: Lustre file system disk storage hardware and 10 GigE federated network switching infrastructure will be enhanced and deployed in preparation for the Sequoia Project's initial delivery and code scalability platform (Dawn ID). This milestone will also provide higher bandwidths to HPSS archival storage in support of Sequoia file transfers.

Milestone (ID#): Performance assessment of the full Roadrunner System
Level: 2
Fiscal Year: FY09
DOE Area/Campaign: ASC
Completion Date:
ASC nWBS Subprogram: Computational Systems & Software Environment
Participating Sites: LANL
Participating Programs/Campaigns: ASC
Description: Assess the performance of the full system RR on a realistic application workload. The work will include analyzing the differences between the achieved/measured performance and the achievable/modeled performance on full applications and optimize the system to maximize performance.

Milestone (ID#): Evaluate advanced memory subsystems (SNL)
Level: 2
Fiscal Year: FY09
DOE Area/Campaign: ASC
Completion Date: Q2
ASC nWBS Subprogram: Computational Systems and Software Environment (CSSE)
Description: Develop, with industry and academia partnerships, an advanced memory subsystem to increase the effective performance of SNL application workloads on next generation microprocessors. Advanced functionality includes (but not limited to) the following operations: atomic memory operations, scatter/gather, in-memory copy/zero/fill/etc, and in-memory synchronization. The architecture will be defined, validated and documented with an industrial partner agreement to develop prototypes.

Milestone (ID#): Deliver smart simulation/experimental petascale data comparison tools (SNL)
Level: 2
Fiscal Year: FY09
DOE Area/Campaign: ASC
Completion Date: Q4
ASC nWBS Subprogram: Computational Systems and Software Environment (CSSE)
Description: We will extend our characterization and comparison tools to handle the sensible comparison of ensembles of large scale simulation data, to supposedly matching experimental data, data that is most likely described and characterized very differently. This is a critical enabling capability to follow through on our programmatic commitments for UQ and V&V. Note that objective, quantitative comparison to experimental data, especially legacy experimental data, is considered a barrier due to technical difficulty and challenges of managing and analyzing petascale data.

Milestone (ID#): Evaluation of the impact chip multiprocessors have on Sandia application performance (SNL)
Level: 2
Fiscal Year: FY09
DOE Area/Campaign: ASC
Completion Date: Q4
ASC nWBS Subprogram: Computational Systems and Software Environment (CSSE)
Description: SNL will investigate the impact of Chip Multi-Processors (CMPs) on the performance of important SNL application codes and the impact of CMPs on the performance and applicability of SNL's system software. This investigation will involve testing of the application codes and system software on several available CMPs in the FY'09 timeframe. A report detailing the results of the investigation will be completed.

Milestone (ID#): Deploy Tripod capabilities for capacity computing environment (TRI)
Level: 2
Fiscal Year: FY09
DOE Area/Campaign: ASC
Completion Date: Q4
ASC nWBS Subprogram: Facility Operations and User Support (FOUS)
Description: Deploy additional Tripod capabilities for capacity computing environment, working towards a responsive and more efficient infrastructure to support computing for QMU and predictivity.

VI. ASC Roadmap Drivers for FY08–FY09

Table VI-1. ASC Roadmap Drivers fro FY08-09[7]

Focus Area 1. Address National Security Simulation Needs

Capabilities for Focus Area 1	Targets for Focus Area 1
1.1 Deliver Simulation Capability for Nuclear Weapons Needs (2007–2009)	2008: National Code Strategy
1.2 Deliver Simulation Technology for Broader National Security Needs (2007–2012)	2009: Modular Physics and Engineering Packages For National Weapons Codes
1.3 Investigate and Understand Off-Normal System Performance and Failure/Anomaly Issues in the Test Database (2008–2013)	

Focus Area 2. Establish a Validated Predictive Capability for Key Physical Phenomena

Capabilities for Focus Area 1	Targets for Focus Area 1
Deliver A Program Plan and Establish A Collaborative Center to Achieve Predictive Physics Capability (2007–2008)	2007: Launch Thermonuclear Burn Initiative Collaboration
Apply Theory, Simulations, Experimental Data, and Interim Model S to Confirm Dominant Physical Phenomena (2008–2009)	2008: Realistic Plutonium Aging Simulations
Demonstrate Improved Physics Understanding Based On Relevant Experimental Results (2009–2013)	2009: Science-Based Replacement for Knob (Ad Hoc Model) #1

Focus Area 3. Quantify and Aggregate Uncertainties in Simulation Tools

Capabilities for Focus Area 3	Targets for Focus Area 3
3.1 Establish and Prioritize the Parameters Matrix (2007–2008)	2008: National Verification & Validation Strategy
3.2 Establish National and Collaborative Forums to Develop Uncertainty Aggregation Methodologies and Benchmarks (2008–2009)	2008: Assessment of Major Simulation Uncertainties
3.3 Deliver Uncertainty Aggregation for QMU Applications (2008–2014)	2009: Shared Weapons Physical Databases

[7] The ASC Top Ten Risks table was originally published in the *ASC Program Plan FY05*.

Focus Area 4. Provide Mission-Responsive Computational Environments

Capabilities for Focus Area 4

4.1 Deploy a Computational Environments For Uncertainty Quantification (UQ) Analyses (2007–2008)

4.2 Deploy Computational Environments and User Facilities for Weapon Science Studies and Other Capability Computing Needs (2008-2012)

Targets for Focus Area 4

2007: Initiate New National User Facility Model for Capability Supercomputing

2008: Seamless User Environments for Capacity Computing

2009: Petascale Computing

VII. ASC Risk Management

Risk management is a process for identifying and analyzing risks, executing mitigation and contingency planning to minimize potential consequences of identified risks, and monitoring and communicating up-to-date information about risk issues. Risk management is about identifying opportunities and avoiding losses. A "risk" is defined as (1) a future event, action, or condition that might prevent the successful execution of strategies or achievement of technical or business objectives, and (2) the risk exposure level, defined by the likelihood or probability that an event, action, or condition will occur, and the consequences, if that event, action, or condition does occur. Table VII-1 summarizes ASC's top ten risks, which are managed and tracked.

Table VII-1. ASC's Top Ten Risks[8]

No	Risk Description	Risk Assessment			Mitigation Approach
		Consequence	Likelihood	Risk Exposure	
1	Compute resources are insufficient to meet capacity and capability needs of designers, analysts, DSW, or other Campaigns.	High	High	HIGH	Integrate program planning with DSW and other Campaigns, to ensure requirements for computing are understood and appropriately set; maintain emphasis on platform strategy as a central element of the program; pursue plans for additional and cost-effective capacity platforms.
2	Designers, analysts, DSW, or other Campaign programs lack confidence in ASC codes or models for application to certification /qualification.	Very High	Low	MEDIUM	Maintain program emphasis on V&V; Integrate program planning with DSW and other Campaign programs to assure requirements needed for certification/qualification are properly set and met.

[8] The ASC Top Ten Risks table was originally published in the *ASC Program Plan FY05*.

No	Risk Description	Risk Assessment			Mitigation Approach
		Consequence	Likelihood	Risk Exposure	
3	Inability to respond effectively with Modeling & Simulation (M&S) capability and expertise in support of stockpile requirements – near or long term, planned or unplanned (SLEP, SFIs, etc.).	Very High	Low	MEDIUM	Integrate program planning, particularly technical investment priority, with DSW and other Campaign programs to ensure capability and expertise is developed in most appropriate areas; retain ability to apply legacy tools, codes, models.
4	Base of personnel with requisite skills, knowledge, and abilities erodes.	High	Low	MEDIUM	Maintain emphasis on "best and brightest" personnel base, with Institutes, Research Foundations, and University programs, as central feeder elements of the program.
5	Advanced material model development more difficult, takes longer than expected.	Moderate	High	MEDIUM	Increase support to physics research; pursue plans for additional computing capability for physics and engineering model development
6	Data not available for input to new physics models or for model validation.	High	Moderate	MEDIUM	Work with Science and Engineering Campaigns to obtain needed data; propose relevant experiments.
7	Infrastructure resources are insufficient to meet designer, analyst, DSW, or other Campaign program needs.	High	Low	MEDIUM	Integrate program planning with DSW and other Campaigns, to ensure requirements for computing are understood and appropriately set; maintain emphasis on system view of infrastructure and PSE strategy, as central elements of the program.
8	External regulatory requirements delay program deliverables by diverting resources to extensive compliance-related activities	Moderate	Low	MEDIUM	Work with external regulatory bodies to assure that they understand NNSA's mission, ASC's mission, and the processes to set and align requirements and deliverables, consistent with applicable regulations.

No	Risk Description	Risk Assessment			Mitigation Approach
		Consequence	Likelihood	Risk Exposure	
9	Inadequate computational environment impedes development and use of advanced applications on ASC platforms.	Moderate	Very Low	LOW	Integrated planning between program elements to anticipate application requirements and prioritize software tools development and implementation.
10	Fundamental flaws discovered in numerical algorithms used in advanced applications require major changes to application development.	Moderate	Very Low	LOW	Anticipate or resolve algorithm issues through technical interactions on algorithm research through the Institutes, ASC Centers, and academia, and focus on test problem comparisons as part of software development process.

VIII. Performance Measures

Table VIII-1. ASC Performance Measures

ADVANCED SIMULATION AND COMPUTING (ASC) CAMPAIGN								
Goal: Provides leading edge, high-end simulation capabilities to meet weapons assessment and certification requirements, including weapon codes, weapon science, platforms, and computer facilities.								
INDICATOR	**ANNUAL TARGETS**							**ENDPOINT TARGET DATE**
	FY04	**FY05**	**FY06**	**FY07**	**FY08**	**FY09**	**FY10**	
Peer-reviewed progress in completing milestones, according to a schedule in the ASC Campaign Program Plan, in the development and implementation of improved models and methods into integrated weapon codes and deployment to their users (long-term output). Panel Criteria: (1) Delivery and implementation of validated models into code projects, and (2) Documented verification of approximations.	High Fidelity Primary Code	Initial baseline Primary Code	Initial validated simulation code for W76 and W80	W80 code baseline	Conduct modern baseline of all enduring stockpile systems	Complete modern baseline of all enduring stockpile systems	Quantify margins and uncertainties of modern baseline simulations	By 2015, accomplish full transition from legacy design codes to modern ASC codes with documented quantification of margins and uncertainties of simulation solutions.
Cumulative percentage of the 31 weapon system	32%	38%	51%	67%	87%	96%	100%	By 2010, analyze 100 percent of 31

ADVANCED SIMULATION AND COMPUTING (ASC) CAMPAIGN								
Goal: Provides leading edge, high-end simulation capabilities to meet weapons assessment and certification requirements, including weapon codes, weapon science, platforms, and computer facilities.								
INDICATOR	**ANNUAL TARGETS**							**ENDPOINT TARGET DATE**
	FY04	**FY05**	**FY06**	**FY07**	**FY08**	**FY09**	**FY10**	
components, primary/secondary/engineering system, analyzed using ASC codes, as part of annual assessments and certifications (long-term output).								weapon system components using ASC codes, as part of annual assessments and certifications (interim target).
The maximum individual platform computing capability delivered, measured in trillions of operations per second (teraflops) (long-term output).	40	100	200	200	200	350	350	BY 2009, deliver a maximum individual platform computing capability of 350 teraFLOPS.
Total capacity of ASC production platforms attained, measured in teraflops, taking into consideration procurements & retirements of systems (long-term output).	75	172	160	360	470	980	980	By 2009, attain a total production platform capacity of 980 teraFLOPS.

ADVANCED SIMULATION AND COMPUTING (ASC) CAMPAIGN								
Goal: Provides leading edge, high-end simulation capabilities to meet weapons assessment and certification requirements, including weapon codes, weapon science, platforms, and computer facilities.								
INDICATOR	**ANNUAL TARGETS**							**ENDPOINT TARGET DATE**
	FY04	**FY05**	**FY06**	**FY07**	**FY08**	**FY09**	**FY10**	
Average cost per teraflops of delivering, operating, and managing all SSP production systems in a given fiscal year (efficiency measure).	$8.15M	$5.7M	$3.99M	$2.79M	$1.96M	$1.37M	$0.96M	By 2010, attain an average cost of $0.96 M per teraflops of delivering, operating, and managing all SSP production systems.

Appendix A. Glossary

APC	Analysis Process Coordinator
ASC	Advanced Simulation and Computing
CAD	Computer Aided Design
CCC	Capability Computing Campaign
CCF	Central Computing Facility
CSSE	Computational Systems and Software Environment (WBS 1.5.4)
CU	Connected Unit
DNT	Defense and Nuclear Technologies Directorate at LLNL
DSW	Directed Stockpile Work
ESN	Enterprise Secure Net
FOUS	Facility Operations and User Support
GA	General Availability
GB	Gigabytes
GB/sec.	Gigabytes per Second
GPFS	Global Parallel File System
HEDP	High Energy Density Physics
HPC	High Performance Computing
HPSS	High Performance Storage System
I/O	Input/Output
ID	Initial Delivery
LAN	Local Area Network
LANL	Los Alamos National Laboratory
LDAP	Lightweight Directory Access Protocol
LDCC	Laboratory Data Communications Center
LLNL	Lawrence Livermore National Laboratory
MPI	Message Passing Interface
NAS	Network-Attached Storage
NFS	Network File System
NIF	National Ignition Facility
NNSA	National Nuclear Security Administration
NPR	Nuclear Posture Review

NSA	National Security Agency
nWBS	National Work Breakdown Structure
OCF	Open Computing Facility
PaScalBB	Parallel Scalable Back Bone Concept
PSI	Parallel Storage Interface
QMU	Quantification of Margins and Uncertainties
R&D	Research and Development
RAS	Reliability, Availability, and Serviceability
RDMA	Remote Direct Memory Access
RRW	Reliable Replacement Warhead
SCC	Nicolas C. Metropolis Center for Modeling and Simulation
SCF	Secure Computing Facility
SFI	Significant Finding Investigation
SLEP	Stockpile Life Extension Program
SNL	Sandia National Laboratories
SSP	Stockpile Stewardship Program
STS	Stockpile-to-Target Sequence
SU	Scalable Unit
TLCC	Tri-Lab Linux Capacity Cluster
TLCC07	Tri-Lab Linux Capacity Cluster for 2007
TOSS	Tripod Operating System Software
TSF	Terascale Simulation Facility
UQ	Uncertainty Quantification
V&V	Verification and Validation
WAN	Wide Area Network
WBS	Work Breakdown Structure
ZFS	Zettabyte File System

Appendix B. Points of Contact

WBS	Title	Contact
1.5.4	Computational Systems and Software Environment	Steve Louis, LLNL, 925-422-1550, stlouis@llnl.gov John Thorp, LANL, 505-665-82265, thorp@lanl.gov Sudip Dosanjh, SNL, 505-845-7018, ssdosan@sandia.gov
1.5.4.1	Capability Systems	Mark Seager, LLNL, 925-423-3141, seager1@llnl.gov Manual Vigil, LANL, 505-665-1960, mbv@lanl.gov Sudip Dosanjh, SNL, 505-845-7018, ssdosan@sandia.gov
1.5.4.2	Capacity Systems	Mark Seager, LLNL, 925-423-3141, seager1@llnl.gov Manual Vigil, LANL, 505-665-1960, mbv@lanl.gov Sudip Dosanjh, SNL, 505-845-7018, ssdosan@sandia.gov
1.5.4.3	Advanced Systems	Mark Seager, LLNL, 925-423-3141, seager1@llnl.gov John Thorp, LANL, 505-665-82265, thorp@lanl.gov Sudip Dosanjh, SNL, 505-845-7018, ssdosan@sandia.gov
1.5.4.4	System Software and Tools	Steve Louis, LLNL, 925-422-1550, stlouis@llnl.gov John Thorp, LANL, 505-665-82265, thorp@lanl.gov Jim Ang, SNL, 505-844-0068, jaang@sandia.gov
1.5.4.5	I/O, Storage Systems, and Networking	Steve Louis, LLNL, 925-422-1550, stlouis@llnl.gov John Thorp, LANL, 505-665-82265, thorp@lanl.gov Mike Hardwick, SNL, 925-294-2157, mfhardw@sandia.gov
1.5.4.6	Pre- and Post-Processing Environments	Steve Louis, LLNL, 925-422-1550, stlouis@llnl.gov John Thorp, LANL, 505-665-82265, thorp@lanl.gov Dino Pavlakos, SNL, 505-844-9089, cjpavla@sandia.gov
1.5.5	Facility Operations and User Support	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Cheryl Wampler, LANL, 505-667-5243, clw@lanl.gov John Zepper, SNL, 505-845-8421, jdzeppe@sandia.gov
1.5.5.1	Facilities, Operations, and Communications	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Steve Shaw, LANL, 505-606-0203, shaw@lanl.gov John Noe, SNL, 505-844-5592, jpnoe@sandia.gov
1.5.5.2	User Support Services	Brian Carnes, LLNL, 925-423-9181, carnes1@llnl.gov Steve Shaw, LANL, 505-606-0203, shaw@lanl.gov Dino Pavlakos, SNL, 505-844-9089, cjpavla@sandia.gov

WBS	Title	Contact
1.5.5.3	Collaborations	Lynn Kissel, LLNL, 925-423-7940, kissel1@llnl.gov Cheryl Wampler, LANL, 505-667-5243, clw@lanl.gov Martin Pilch, SNL, 505-845-3047, mpilch@sandia.gov
1.5.1.4.12	Caltech, Center for Simulating Dynamic Response of Materials	Dan Meiron, 626-395-3424, dim@its.caltech.edu Michael Aivazis, 626-395-3424, aivazis@caltech.edu Paul Dimotakis, 626-395-6811 Bill Goddard, 626-395-2731, wag@wag.caltech.edu Michael Ortiz, 626-395-4530, ortiz@aeor.caltech.edu Dale Pullin, 626-395-6081, dale@galcit.caltech.edu
1.5.1.4.12	Stanford, Center for Integrated Turbulence Simulation	Parviz Moin, 650-723-9713, moin@stanford.edu Juan Alonso, 650-723-9954, jjalonso@stanford.edu Heinz Pitsch, 650-736-1995, h.pitsch@stanford.edu William Dally, 650-725-8945, billd@csl.stanford.edu Pat Hanrahan, 650-723-8530, hanrahan@cs.stanford.edu
1.5.1.4.12	University of Chicago, Center for Astrophysical Flash Phenomena	Donald Q. Lamb, 773-702-7194, d-lamb@uchicago.edu Robert Rosner, 773-702-0560, r-rosner@uchicago.edu Todd Dupont, 773-702-3485, Dupont@cs.uchicago.edu Rusty Lusk, 630-252-7852, lusk@mcs.anl.gov Tomasz Plewa, 773-834-3227, tomek@flash.uchicago.edu Andrew Siegel, 773-834-8501, siegela@flash.uchicago.edu Jim Truran, 773-702-9584, truran@nova.uchicago.edu
1.5.1.4.12	University of Illinois, Center for Simulation of Advanced Rockets	Michael T. Heath, 217-333-6268, m-heath@uiuc.edu William A. Dick, 217, 244-7235, wdick@uiuc.edu Robert A. Fiedler, 217-333-3247, rfiedler@uiuc.edu
1.5.1.4.12	University of Utah, Center for Simulation of Accidental Fires and Explosions	David W. Pershing, 801-581-5057, David.Pershing@utah.edu Chuck Wight, 801-581-8796, Chuck.Wight@utah.edu Tom Henderson, 801-581-3601, tch@cs.utah.edu Philip Smith, 801-585-3129, smith@crsim.utah.edu Patrick McMurtry, 801-581-3889, mcmurtry@eng.utah.edu Eric Eddings, 801-585-3931, eddings@che.utah.edu Steve Parker, 801-585-1504, sparker@cs.utah.edu

Appendix C.
WBS 1.5.1.4-TRI-001 Alliance Support

For FY08, the Academic Strategic Alliance Program (ASAP) Centers will be performing work on a six-month extension to their recently completed ten-year contracts. A summary of their planned work follows.

Stanford Center for Integrated Turbulence Studies

Simulation Plan and Issues to Be Studied

The current simulations are based on a 20-degree sector of the full engine and are performed using a relatively coarse mesh. The plan for the extension is to perform a simulation on a refined grid (consisting of 80 million elements). Additional grid convergence studies will be performed on the engine components (not integrated) by considering finer resolution. A study using a calculation in a 60-degree sector is also being considered to alleviate some concerns related to geometric scaling of the turbo machinery due to variable blade counts in different sectors. These simulations will be performed depending on the availability of Red Storm. In addition

Additional multi-physics models, such as primary fuel break-up and soot formation will be incorporated in the end-end simulations.

The key objective of the end-end engine simulations is to study the effects of component interactions on the full engine performance. In particular, studies will be made of effects of the fan blade wakes on the low pressure compressor and the migration of hot-streaks from the combustor into the turbine.

Verification

A key area of the verification effort will be establishing grid convergence of large-eddy simulations by filtering of the governing equations. LES simulations will be carried out with explicit filtering to demonstrate grid convergence of the numerical solutions. Another area of study will be verification of multi-physics problems that involve various physical models. Focus will be on developing manufactured solutions for two aspects of reactive flow simulations – fuel droplet tracking and combustion modeling.

Validation

Close collaboration with P&W on the analysis of the Center results using proprietary data will continue. Validation experiments have been commissioned at Stanford, Michigan State and Notre Dame. Simulations and experiments of acoustics and turbulence in airfoil wakes will be made. Additional validation studies will be performed.

Uncertainty Quantification

Focus will be on identifying uncertainties at code interfaces and to assess their impact on predictions. Also uncertainties associated with the prediction of flow and heat transfer in the turbine will be assessed.

Parallel Programming Systems

The Sequoia stream-based parallel programming model/language and compile/runtime will continue to be developed, ported to Roadrunner at LANL (and considered for a port to GPUs). The implementation of SUmb in Sequoia will also be extended.

University of Illinois Center for the Study of Advanced Rockets

Simulation Plan

The primary simulation target will remain the Reusable Solid Rocket Motor (RSRM) that powers the U.S. Space Shuttle. During the six-month extension period, the goal will be to simulate the RSRM with improved physical fidelity so that they can explore a number of important science and engineering issues (see below) and validate the results against experimental data whenever possible.

Detailed Simulation Studies

Nozzle Ablation: Assessing the degree of possible ablation or erosion of the rocket nozzle due to particle impingement is of critical interest. Multiphase flow statistics at the nozzle entrance inlet will be extracted from *Rocstar* simulations of several critical physical phenomena and parameters. Simulation predictions will be compared with experimental data, where available, to validate insights gained into the mechanisms of nozzle ablation.

Multiphase Plume Dynamics: Exhaust plume dynamics can have a significant effect on surrounding structures in complex applications such as the Space Shuttle. *Rocstar* simulations have been used to study exhaust plume dynamics for lab-scale motors. The computational domain will be extended to include a large plume region in which results will be analyzed for the RSRM.

Slag Accumulation: Solid propellants are often enriched with aluminum particles to increase specific impulse and damp combustion instabilities, but the resulting accumulation of slag (liquid aluminum oxide), for example pooling in a submerged nozzle, can lead to a drag on performance and potential sloshing or ejection that can cause control problems or thrust spikes. Slag accumulation in large solid rocket motors will be simulated for a variety of conditions.

Acoustics in Solid Rocket Motors: Understanding all sources that contribute to the acoustics in solid rocket motors remains an unsolved problem in rocket design. As a result, resonances can unexpectedly occur in a motor leading to catastrophic failure. Based on simulations of heterogeneous propellants using *Rocfire,* there is evidence that any simulations ignoring the time scales associated with unsteady burning are fundamentally flawed. Thus this critical area will be studied.

Effect of Heterogeneity on Turbulence: The spatial and temporal fluctuations in the flow field arising from heterogeneity will be examined using fully 3-D simulations for the BATES and StarAft motors using a variety of propellant morphologies to study their impact on performance.

Additional Verification and Validation Activities

Advanced Imaging Technologies for Validation: Substantial effort has been devoted over the past two years to the characterization of solid propellant microstructure, unit cell reconstruction, and the generation of computational domains (meshing). Micro-tomography will be used to characterize the solid propellant and validate the packing algorithm.

As part of an ongoing effort to gather appropriate data for further validation of the multiscale and multiphysics capabilities of *Rocstar*, funds from DOE have contributed to the purchase of a 2 billion frame per second camera. The camera will be used for investigations of multiscale dynamic properties in heterogeneous energetic materials.

Verification of Time Zooming: The propellant of the ACM rocket can be burned out in a week or two of simulation wall clock time without time zooming, which makes it a good verification case for both the structured and unstructured fluid solvers. Since pressure history data are available, this motor has been used extensively in UQ studies. Particularly interest is focused on behavior close to burnout. Plans are also underway to use the StarAft motor for verifying time zooming.

Verification of Staggered Coupling: The staggered coupling scheme used in *Rocstar* will be further verified by comparing results for model fluid-structure interaction problems with results produced by a monolithic fluid-structure interaction code.

Sample-Based UQ: Considerable progress has been made in the past year on sample-based surrogate-model clustered uncertainty propagation through rocket simulation models. For the proposed extension period, the more mature surrogate-clustering UQ techniques will be applied to the full-scale RSRM, with the goal of generating a time-pressure envelope to compare with the variability that is evident in available RSRM firing data.

Additional Science and Engineering Developments

A number of additional simulation capabilities will be developed and incorporated into both the basic physics solvers and the integration framework underlying *Rocstar* to support the simulations described above

- Implement a non-dissipative, fully implicit all-speed algorithm in *Rocflu* to improve its capabilities in predicting turbulent flows.
- Implement time zooming in the block structured fluid solver *Rocflo*.
- Enhance meshing support to make *Rocstar* more robust for full burns and large geometric changes.
- Implement fully parallel remeshing, parallel mesh repair.
- Incorporate our fully parallel, hybrid geometric/topological mesh partitioner into *Rocstar* and integrate it with remeshing.
- Continue enhancements to our particle packing algorithm.
- Extend our techniques for load balancing, fault tolerance, and performance prediction to petascale machines.

University of Chicago Center for the Study of Thermonuclear Flashes

During the six-month extension period, the Flash Center will address the following three significant scientific questions:

- Does Rayleigh-Taylor (R-T)-driven turbulent nuclear burning occur mostly at small or at large scales?
- Can the gravitation-ally confined detonation (GCD) mechanism explain the full range of peak luminosities of Type Ia supernovae?
- What are the properties of the light curves, spectra, and polarizations predicted by 3-D simulations of the GCD model?

Simulations Plan

The following simulations will be conducted during the extension period.

- 3D AFL simulations for planar and bubble initial conditions. A sequence of 642 x 256, 128 x 512, and 256 x 1024 simulations on uP are underway. These will lead up to several 512 x 2048 simulations on Argonne BG/P in the fall, each of which will take - 0.5M cpu-hrs.
- 3D simulations of the transition from the flamelet nuclear burning regime to the distributed nuclear burning regime due to fully developed R-T turbulence. A sequence of 642 x 256, 128 x 512, and 256 x 1024 simulations on uP are underway. These will lead up to a 512 x 2048 simulation (and possibly a 1024 x 4096 simulation) on Argonne BG/P in late 2007; the first will take - 0.5M cpu-hrs and the second 8M cpu-hrs.
- First exploratory 3D simulations of the Wisconsin shock-tube bubble experiments. The initial focus will be on getting the initial conditions correct. The first bubble simulations will be done in late 2007.
- 3D simulations of GCD mechanism for multiple off-center ignition points. Each simulation will take - 0.1M cpu-hrs. The treatment of gravity and the algorithm for creating the initial model are ready. Verification simulations are underway.
- 3-D simulations of the GCD mechanism in the presence of convective turbulence in the core of the white dwarf. Each simulation will take 0.1M Pu-hrs. We will first explore the effect of a turbulent velocity field on bubble evolution, using AFL simulations. We will then implement the same initial conditions in 3-D simulations of the GCD model.
- 3D radiation transfer simulations of light curves and spectra predicted by the GCD mechanism using 3-D hydro results and nuclesynethetic yields from post-processing of Langrangian tracer particles. A set of 2-D pathfinder radiation transfer simulations are complete and have been published. Initial 3-D radiation transfer simulations will begin in the fall and continue in Winter 2008.

Verification & Validation

The following additional verification and validation (V&V) activities will be conducted to support the simulations that will be done during the extension period: (1) Completion of 3D simulations of the Michigan shock-tube R-M/R-T experiments; (2) initiation of simulations of the Wisconsin shock tube bubble experiments and carrying out of the first reactive bubble experiments; and (3) application of the Gaussian Process methodology that we have developed to data from 2D and 3D simulations of the light curves and spectra predicted by the GCD model and from the SDSS-II Supernova Survey.

The plan for interacting with the observational community during the extension period includes (1) biweekly SNacks with members of the SDSS-II Supernova Survey team; (2) monthly telecons with members of the SDSS-II, Harvard, and Magellan Supernova Survey teams; (3) a series of visits to the Flash Center by members of these teams; (4) a second annual Type Ia Supernova Workshop to be held in Chicago in September 2007; and (5) informal interactions at international conferences and workshops.

Science and Engineering Issues

The following additional developments will be needed to support the simulations that will be carried out during the extension period:

- Implementation of algorithm for constructing initial white dwarf model with convective core for 3-D GCD simulations;
- Implementation of version of Dan Karen's *Sedona* Monte Carlo radiation transfer code optimized for massively parallel computers;
- Migration of FLASH 3 to Red Storm; and
- Migration of FLASH 3 to Argonne BG/P, which will be funded by the DOE Office of Science.
- An MPI profiling library for detecting, at runtime, consistency errors in the use of MPI collective operations;
- Thread-safe, lock-free logging of MPI calls and other program events and display of traces of multi-threaded programs; and
- Development of a dynamic load-balancing library for massively parallel, multi-core machines.

Visualization developments that will be needed are tools and infrastructure to provide rapid feedback to the scientists, and enhanced capabilities to generate movies and images needed for visual analysis of the large data sets now being generated. Specifically,

- Additions to the LLNL Visit tool that are needed by the Flash Center
- Continued development of native volume rendering toolkit for FLASH data
- Development of an analysis pipeline and data archive.

University of Utah Center for the Simulation of Accidental Fires and Explosions (C-SAFE)

Since 1997, the University of Utah Center for Simulation of Accidental Fires and Explosions (C-SAFE) has been building a state-of-the-art science-based high-performance simulation code to predict the explosive response of a steel container filled with PBX9501 embedded in a large jet fuel fire.

Oxygen-Fired Coal Burner Simulations

This topic was chosen because of the huge potential impact that it can have in the effort to retrofit existing air fired coal burners for the purpose of sequestering the carbon dioxide product gases. Combustion of coal in air is a major contributor to greenhouse gas emissions, and the rise of China, India and other industrialized countries threatens to exacerbate the effects of global warming into the foreseeable future. In the long term, two new technologies (Integrated Gasification Combined Cycle (IGCC) and Ultra Supercritical Oxy-Fuel boilers) have the potential to provide high efficiency power plants capable of near zero pollutant emissions with integrated CO_2 capture. However, these technologies will require extensive development cycles and huge capital investment. In addition, neither are suitable for application to the more than 500 existing coal-fired power plants in the US alone.

Promising shorter-term solutions include post-combustion amine scrubbing and oxygen-fired coal burner retrofits. The principal objective during this six-month extension period is to demonstrate the feasibility of firing coal in mixtures of oxygen and recycled carbon dioxide by using LES to predict flame stability, emissions, and heat transfer. This will enable C-SAFE to then partner with burner manufacturers to design burners that can be used in existing boilers, thereby achieving maximum potential for

reducing carbon emissions at minimum capital cost. These simulations could also be valuable for helping analyze future designs, such as Ultra Supercritical Oxy-Fuel boilers.

Required Simulations: Using ARCHES, three LES coal scenarios will be demonstrated with computational domains focused in the near-burner region. These scenarios are 1) a cold-flow co-annular jet; 2) an oxy-fuel fired, co-annular jet; and 3) an air-fuel fired, co-annular jet. Each simulation will use roughly 80,000 CPU hours to determine the effects of the oxygen content on flame stability, emissions, and heat transfer in the boiler. These scenarios will require multiple simulations to investigate model parameters, grid sensitivity, and validation studies. The simulations will require a successful port of Uintah and supporting solver codes to the Sandia Red Storm machine.

Additional Science/Engineering Advancements: The ARCHES code within Uintah does not currently have the ability to handle a solid disperse phase, an essential feature for performing oxy-coal simulations.

Verification and Validation: Using experimental data obtained by J. Sinclair Curtis et al. (Purdue) for non-reacting particle laden jets and Eddings et al. for oxy-fuel coal jets, simulations performed in the three scenarios will be validated. Additional verification and validation tests will be conducted for submodels to provide the desired level of confidence in the simulated results.

Explosion Violence and Device Geometry

As part of the ongoing validation work in C-SAFE, Eddings et al. demonstrated that the violence of explosion is more sensitive that previously believed to both the heating rate and the presence of a hollow bore. Uintah simulations showed the most likely cause of violent explosions at moderate to slow heating rates is thermal penetration of the PBX9501 inside a 4-inch diameter steel container, which causes the collapse of the hollow bore following ignition. Work performed under the 6-month contract extension will be directed toward detailing the dependence of violence on device geometry.

Required Simulations: Scoping calculations in 2D have been completed, so the main focus will be on full 3D simulations to quantify reaction violence as a function of bore size for a 4-inch ID steel container filled with PBX9501.

Additional Science/Engineering Advancements: These simulations will utilize existing AMR strategies for focusing the computational power on the solid device undergoing combustion and explosion. The 3D simulations will take advantage of code efficiency improvements currently underway.

Verification and Validation: Fully instrumented full-scale validation experiments have already been completed for this project.

Sensitivity Analysis

In year 10, C-SAFE performed a sensitivity analysis on the experimental scenario of a 10 cm-diameter container of energetic material (10 cm in length) being heated in a JP-8 pool fire. The objective was to investigate sensitivity of explosion violence in end-to-end simulations on a range of heating rate conditions generated by varying three variables: pool fire diameter, wind speed, and position of the container relative to the pool.

C-SAFE personnel will extend this sensitivity analysis to include critical sub-scale computational models. These analyses will focus on the dependence of end-to-end simulation predictions on the parameters of the submodels and on the scientific content of the models.

The Uintah simulation code incorporates either of two state-of-the-art PBX combustion models. The first assumes that combustion is always at steady state and depends only

on local temperatures (bulk PBX and surface temperatures) and local pressure (which determines the flame standoff distance). There are 12 parameters that determine the behavior of this steady combustion model. A sensitivity analysis using Monte Carlo Latin Hypercube selection of parameter values will be conducted in order to determine the sensitivity of burn rates to the various parameters.

The sensitivity analysis will include a determination of the dependence of combustion rate on inclusion of the additional physics of unsteady combustion, even under conditions where steady combustion rates computed by the two models are identical.

The fully molecular EOS/MM, which includes effects of temperature and pressure, will be compared to the ViscoSCRAM model from LANL, which includes only static parameters but has been the gold standard for PBX9501 up to now.

On the fire side, sensitivity analyses will focus on resolved-scale and subgrid-scale models that significantly affect heat transfer to containers immersed in or adjacent to JP-8 pool fires. Sensitivity of fire/container heat transfer to the choice of soot model will be tested . Sensitivity to the fuel burning rate will also be tested with the implementation of a fuel evaporation model that is based on heat transfer from the fire to the pool surface.

C-SAFE will also conduct a sensitivity analysis of the integrated end-to-end code by determining the sensitivity of the predictions (time to explosion and explosion violence) to changes in models and input parameters.

Required Simulations: Some of the sensitivity analyses described in this section can be performed in 2D on modest numbers of processors. However, 3D simulations involving PBX combustion model sensitivity, fire sensitivity to soot models and to fuel burning rate, and sensitivity of the integrated end-to-end code to various models will require 12-15 full-scale runs on National Laboratory machines. Based on experience, each run will require approximately 50,000 CPU-hr.

Additional Science/Engineering Advancements: Sensitivity analyses of burn models and the end-to-end code can be done with the current Uintah code. Comparison of the material models requires completion of the first-principles EOS/MM and full implementation in the Uintah computational framework. Scoping studies have shown that there are no difficult obstacles to this, and completion of the initial version of the EOS/MM is anticipated by early Fall 2007. Improvements to radiative heat transfer computational algorithms will reduce the computational load and improve the scaling and accuracy of the heat transfer in fire simulations. Currently, radiation consumes over 60% of the computational time in a fire simulation, so efforts will focus on reducing that computational load.

Verification and Validation: Verification and validation of the first-principles EOS/MM will take place during the performance period, including comparisons with known physical properties of PBX and with wave propagation and dynamic response characteristics derived from the Sandia Z-Pinch experiments on explosives. Validation of MPM/ARCHES simulations of heat flux to containers in jet fuel fires will also continue through the extension period. We have identified three relevant data sets: (1) a 1.2 m-diameter steel calorimeter (4.6 m in length) immersed in a 7.2 m-diameter JP-8 pool fire from Kramer et al., 2003, (2) a well-characterized 7.9 m-diameter JP-8 pool fire with equipped with four calorimeters, including two 30-cm diameter calorimeters placed near the pool center and one 1.2 m-diameter calorimeter located adjacent to the pool from Blanchat et al., 2006, and (3) a 3.7 m-diameter steel calorimeter (9.1 m in length) located at the lee end of 9.5 m-diameter and 18.9 m-diameter JP-8 pool fires from Suo-Anttila and Gritzo, 2001. Simulations of the first data set are nearly complete, so validation efforts will focus on cases (2) and (3). A new metric proposed by Oberkampf and Barone (2006) that is based on the statistical concept of confidence intervals is being used.

Caltech Center for the Study of the Dynamic Response of Materials

Simulation Plan

In Phase 2 (FY03-FY07) of the Alliance program the Caltech center focused on the application of the Virtual Test Facility (VTF) to a series of integrated multiphysics simulations each with direct relevance to fundamental scientific issues in the dynamic response of materials that in turn are directly connected to validating experiments. In this six month extension the approach taken in Phase 2 where validation is intimately tied in with simulation will be continued with specific emphasis on several high priority questions.

Plastic deformation due to detonation and shock waves: Work on the deformation of metal tubes under the action of shock and detonation waves will continue. Experimental studies will be used to define the important physical effects and provide test cases for the validation of the fluid-solid structure computations with the VTF. The goal is to obtain larger deformations and examine the possibility of rupture due to plastic instability.

Detonation induced rupture in metal tubes: The investigation of detonation induced ductile rupture of aluminum tubes will continue. The focus will be to understand and remove the significant discrepancy in crack speed predictions of the VTF and the experiments.

Ballistic impact of Homalite 100 target by Ta projectile at velocities in the range of 2-3 Km/s.: The calculations will be based on multiscale models of Ta, adiabatic heating, compressive damage of brittle materials, and discrete cracking in the tensile region, developed previously. The corresponding validation experiments will be carried out in Prof. Ares Rosakis' ballistic impact facility. Validation data will include fragment distribution measurements on witness plates. Validation will be based on a quantitative measure of the distance between the computed and measured fragment distribution functions. The adiabatic heating component of the Ta model will also be validated in the shear-compression specimen configuration using data from Prof. Ravichandran's dynamic testing laboratory

Plate-impact Ta shock experiment at velocities in the range of 2-3 Km/s.: The calculations will be based on multiscale models of Ta, adiabatic heating, direct simulation of polycrystals and explicit constructions of optimal sub-grain microstructures. Both random and rolling textures will be analyzed. Full 3D calculations will be carried out in the parametric regions of greatest interest. Mesh convergence will be carefully verified. Validation will be based on Visar data from plate-impact experiments conducted at Sandia National Laboratories.

Science and Engineering Issues

Cylindrical Richtmyer-Meshkov instability: Several three-dimensional exploratory simulations of Richtmyer-Meshkov instability have been completed. In FY08 six to eight production runs on cylindrical RM instability spanning several parameters ranges including Mach number, heavy-light and light-heavy interfaces are under consideration. The effect of these configurations on the mixing properties will be investigated. The influence of converging geometry on the shock-induced mixing will be studied in relation to the planar case.

Turbulence modeling with realistic EOS: The two-fluid version of the Miller-Puckett algorithm for fluids obeying the Mie Gruneisen equation of state will be implemented into AMROC.

Electron FF: Development of the EFF approach to describing warm dense matter with substantial numbers of molecular species, atoms, ions, and electrons will be completed.

Quantum Monte Carlo: Development of the Caltech Quantum Monte Carlo code will be completed and tested in the Blue Gene type environment and fully implemented on GPUs.

New finite element bases: Construction of FE bases over tet/oct meshes will be continued.

Meshing: More meshing algorithms, for arbitrary 2D and 3D domains, using simplicial or cell complexes will be developed.

New approaches to simulating free surfaces: Study of "foliation processing" approaches to numerical interface handling will be continued.

Fault tolerance for parallel calculations: The group will test the scalability of its fault tolerance framework and its robustness.

Advances in the Pyre framework: Work to enable the creation, deployment and monitoring of a distributed machine entirely within Pyre will continue.

Distribution of ReaxFF: The portability of ReaxFF will be extended and improved for outside distribution. The most important goal will be to consolidate the gains in the development of ReaxFF to make it stable and useful in outside laboratories,

Insertion of high performance computing culture into the curriculum: Continued effort will be made to insert high performance computation into the Caltech curriculum and develop the CSE program.

Software distribution for Pyre and the VTF: The center will continue its implementation of release procedures for Pyre and VTF.

Further Validation

Fragmentation of Homalite targets subjected to impact by Ta projectiles: During the past four years of the validation group's research efforts the team has concentrated on issues relating to the validation of the numerical dynamic fracture capabilities for brittle materials. For the next six months of the extended effort they will use a new experimental set up which will provide a continuation of previous efforts in to the, not yet validated, regime of dynamic fragmentation.

Converging shock experiments: The six-month extension will permit the completion of phase 1 of the converging shock experiment.

Phase transformation of iron under shear compression at finite temperatures: In the proposed 6 month period, a systematic series of experiments will be performed to further explore the mechanics and mechanisms of possible shear induced phase transformation in pure a-iron as a function of temperature and prestrain.

Quantification of discrepancies between numerical simulation and experiment: This will continue quantify discrepancies between numerical simulation and experiment during the six month extension.

Appendix D.
ASC Obligation/Cost Plan

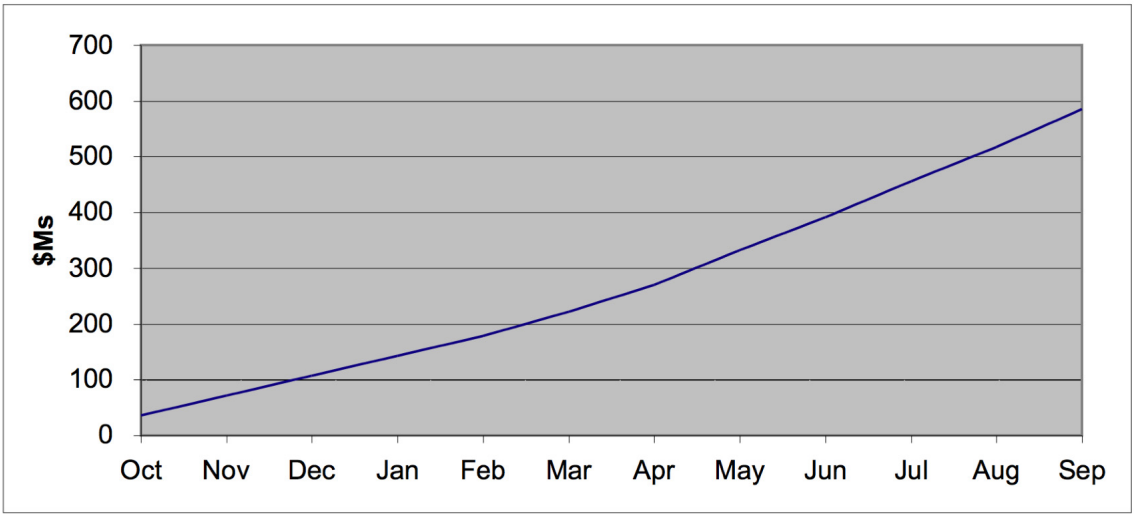

Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep
35.523	71.548	107.472	142.233	177.871	222.056	269.830	331.930	391.497	455.713	517.286	585.738

Figure D-1. ASC obligation/cost plan for FY08.

Appendix E.
ASC Performance Measurement Data for FY08

Table E-1. ASC Performance Measurement Data for FY08

ADVANCED SIMULATION AND COMPUTING (ASC) CAMPAIGN								
Goal: Provide the computational science and computer simulation tools necessary for understanding various behaviors and effects of nuclear weapons for responsive application to a diverse stockpile and scenarios of national security.								
INDICATOR	**ANNUAL TARGETS**							**ENDPOINT TARGET DATE**
	FY07	**FY08**	**FY09**	**FY10**	**FY11**	**FY12**	**FY13**	
CODE PREDICTIVITY: Biennial progress toward a desired end-state through a series of incremental targets of code usability and reliability for ASC applications as measured by the Code Maturity Index.	Baseline	TBD	TBD	TBD	TBD	TBD	TBD	By 2015, achieve 100% of target code maturity as assessed using the ASC Code Predictivity Characteristic Matrix and measured by the Code Maturity Index.
STOCKPILE IMPACT: The cumulative percentage of high consequence stockpile activities to which the modern ASC codes are the predominant simulation tools as measured by the national Stockpile Application Index (nSAI).	Baseline	55%	65%	77%	88%	100%	100%	By 2012, designers and analysts will use ASC codes predominately in 100% of the defined list of high consequence stockpile activities as measured by the nSAI.

ADVANCED SIMULATION AND COMPUTING (ASC) CAMPAIGN								
Goal: Provide the computational science and computer simulation tools necessary for understanding various behaviors and effects of nuclear weapons for responsive application to a diverse stockpile and scenarios of national security.								
INDICATOR	**ANNUAL TARGETS**							**ENDPOINT TARGET DATE**
	FY07	**FY08**	**FY09**	**FY10**	**FY11**	**FY12**	**FY13**	
CAPABILITY COMPUTING: The annual percentage of usage of ASC capability platforms for simulations that use at least 30% of the platform capability, as defined by total available node hours.	Baseline	18%	25%	32%	39%	45%	45%	By 2012, use 45% of the total each ASC capability platform for simulations that use at least 30% of the platform capability.
COST EFFICIENCY: The annual cost ($Ms) per petaFLOPS (quadrillion floating point operations per second) to procure, upgrade, operate, manage and maintain ASC computing platforms.	Baseline	1,089	762	534	320	192	121	By 2015, achieve a decrease in the life-cycle cost per petaFLOPS to $100M from FY2007 Baseline.
Code Predictivity milestones	Deliver a physics-based sub-grid model to support energy balance resolution (LLNL) Enhanced Pu multiphase Equation of State capability (LANL) Predictive failure capabilities in SIERRA mechanics (SNL) Fundamental simulations of material response and plasma physics on HPC platforms (LANL)							
Stockpile Impact milestones	ASC simulations supporting the National Technical Nuclear Forensics Attribution program (LANL) Assessment of nuclear physics uncertainties (LLNL) Aged temperature/radiation aware model of Complementary Metal Oxide Semiconductor integrated circuit technology (SNL)							
Capability Computing milestones	Infrastructure Deployment Plan for ASC Petascale Environments (LLNL, LANL, SNL)							

Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, CA 9455-0808